Revisiting Global Poverty Reduction: Public Services and the World Distribution of Income, 1980-2022

Amory Gethin*

15th July, 2025

Abstract

Global poverty and inequality statistics focus on consumption or disposable income, ignoring transfers received by households in the form of public services. This article combines newly assembled microdata, public spending series, and quality indicators to construct government redistribution measures that include education and healthcare in 173 countries. There have been considerable improvements in public services received by the global poor. From 1980 to 2022, the value of public education and healthcare received by the world's poorest 20% doubled as a share of global GDP. The consumption of public services can account for 30% of growth among the world's poorest 20% since 1980. Total redistribution, including cash and in-kind transfers, can account for 50%. The growth in public education and healthcare services can rationalize the totality of the growing gap between survey and national accounts aggregates documented in the literature.

JEL: D31, H50, I14, I24, I32.

*World Bank Development Research Group; amory.gethin@gmail.com. I am particularly grateful to Thomas Piketty, Vincent Pons, Emmanuel Saez, and Gabriel Zucman for their continuous guidance. I thank Daron Acemoglu, Lars Harhoff Andersen, Alan Auerbach, Pierre Bachas, Thomas Blanchet, Denis Cogneau, Léo Czajka, Benoît Decerf, Edgard Dewitte, Matthew Fisher-Post, Ignacio Flores, Jonathan Gruber, Anders Jensen, Lawrence Katz, Christoph Lakner, Wouter Leenders, Ségal Le Guern Herry, Rosanne Logeart, Clément Mazet-Sonilhac, Marc Morgan, Rowaida Moshrif, Oda Nedregård, Paul Novosad, Tom Raster, Anne-Sophie Robilliard, Justin Sandefur, Todd Schoellman, Camille Semelet, Andrei Shleifer, Morten Nyborg Støstad, David Thesmar, Tom Vogl, Torsten Figueiredo Walter, Duncan Webb, and Ekaterina Zhuravskaya, as well as seminar and conference participants at the Paris School of Economics, UC Berkeley, the University of Namur, and the University of Washington for useful comments and discussions.

1. Introduction

Government redistribution has been on the rise. From 1980 to 2022, worldwide public expenditure per capita doubled in real value. Cash transfers cannot be held sole responsible: they represent less than 10% of public expenditure. Instead, growing redistribution was driven by investments in education, healthcare, housing, police services, transport infrastructure, and other public goods. Together, these transfers represented 30% of global GDP in 2022.[1]

This transformation remains absent from poverty and inequality statistics. The standard concept used to measure global poverty is household final consumption expenditure, defined as the market value of all goods and services purchased by households. By construction, it excludes public services, since these services are not bought on a market. As a result, it remains difficult to understand how macroeconomic growth reduces poverty, in a world where almost a third of global GDP is redistributed by governments in unaccounted ways.

This paper studies redistribution in the form of public education and healthcare and its role in the historical reduction of global poverty. The starting point is a new database on the world distribution of public services from 1980 to 2022. I construct this database by combining data on two dimensions: the total value of public services and their distributional incidence.

The first step is to value public education and healthcare. I consider three alternative methods. My benchmark specification values these services at their cost of provision, drawing on a new historical database on public education and healthcare spending, which is the most common approach in the inequality literature (e.g., Auten and Splinter, 2024; Piketty, Saez, and Zucman, 2018). The main advantage is that it ensures that inequality measures are consistent with the national accounts. The drawback is that public spending may not adequately reflect the quality of services and willingness to pay for them. I thus complement my analysis with two alternative measures. The second measure corrects cost of provision downwards in countries with low

[1] See Appendix Figure A2, which plots the evolution of real worldwide government expenditure per capita.

government productivity, evaluated by comparing education and health outcomes obtained at a given level of spending. The third measure is the net present value of education and healthcare, proxied by expected returns to schooling and gains in life expectancy enabled by the education and healthcare systems. While more challenging to estimate, this measure gets us closer to the welfare value of public services.

The second step is to estimate who benefits from these services. I construct a new microdatabase on inequality in access to education and healthcare. A first set of surveys cover household income and school attendance of children, allowing me to estimate how the consumption of education services varies by income group. A second set of surveys report public healthcare use intensity, proxied by household members' visits to healthcare providers. I complement these indicators with additional data on social assistance transfers (cash transfers and in-kind social benefits such as food stamps), allowing me to compare the incidence of public services to that of these traditional redistributive tools. I also extensively validate my estimates of redistribution by comparing them to existing high-quality studies covering a more restricted number of countries.

The analysis of this new database yields four sets of results.

First, I document large cross-country variations in public education and healthcare provision. Low-income countries spend less on education, healthcare, and social assistance than high-income countries; they provide each of them more unequally; and they deliver them slightly less efficiently. This "triple curse" of redistribution in low-income countries implies large inequalities in the quality of public services received worldwide. One implication is that the share of the global poor living in poor countries is greater than we thought, because poor countries benefit from public services of much lower quality than those of the rich world.

Second, there have been large improvements in the quality of education and healthcare received by the global poor. From 1980 to 2022, the consumption of public services by the world's poorest 20% quadrupled in real value and doubled as a share of global GDP. As a result, ignoring public services would lead to underestimating progress made by tax-and-transfer systems in the reduction of poverty and inequality by 50%. This result holds across all three alternative measures of the value of public services, despite the different assumptions they rely on.

Third, the consumption of public services has played a major role in global poverty and inequality reduction (see Figure 1). The real average income of the world's poorest 20% increased by about 80% in terms of pretax income, 125% after adding social assistance transfers, and 175% after also accounting for the consumption of public services. Total government redistribution thus accounts for about 50% of global bottom 20% income growth during this period. Public education and healthcare alone account for 30%.

Finally, I show that public services can account for all of the growing discrepancy between surveys and national accounts aggregates documented in the literature. Whether it is surveys or GDP that should be used to track global poverty has been an ongoing debate (e.g., Deaton, 2005; Sala-i-Martin, 2006). In an important contribution, Pinkovskiy and Sala-i-Martin (2016) find that nighttime lights correlate much better with GDP than with survey aggregates, implying that GDP should be prioritized. Their analysis reveals another puzzling fact: the gap between GDP and survey means has been growing, and disproportionately so in countries that have seen greater improvements in several indicators of quality of life. I show that this growing gap can be entirely rationalized by progress made in the provision of public education and healthcare, whose characteristic is precisely that they are included in GDP but absent from surveys.

A growing literature has made progress in bridging gaps between micro- and macro-approaches to the measurement of living standards. Piketty, Saez, and Zucman (2018) construct Distributional National Accounts (DINA) for the United States, allocating the entirety of national income to individuals. A number of studies following a comparable methodology have been conducted since then. The advantage of this methodology is that it produces estimates of inequality that are consistent with macroeconomic growth. However, major uncertainties remain when it comes to public services, which are often allocated using *ad hoc* assumptions such as proportionally to income, as a lump sum transfer, or as a combination of both (Auten and Splinter, 2024;

Blanchet, Chancel, and Gethin, 2022; Piketty, Saez, and Zucman, 2018). A few studies have used detailed administrative data in specific contexts, with the general conclusion that public education and healthcare strongly reduce inequality (Bruil et al., 2022; Germain et al., 2021; Gethin, 2024). This paper extends the scope of the analysis by constructing a new database on the world distribution of public education and healthcare. My results also speak to the growing literature on inequality and government redistribution in the developing world (Asher et al., 2021; Bachas, Gadenne, and Jensen, 2024; Jensen, 2022; Lustig, 2018).

This paper also relates to the literature on the valuation of in-kind transfers. A large literature studies the distributional incidence of government expenditure.[2] An important concern is that public spending might not be an adequate measure of the value of public services in the presence of government inefficiencies or unobserved technological progress (e.g., Chong et al., 2014; Cutler et al., 2022; Muralidharan and Sundararaman, 2015). Cost of provision may also differ from willingness to pay, which depends on benefits that vary across policies (e.g., Hendren and Sprung-Keyser, 2020). This paper takes a global historical perspective on these issues, and investigates the implications of addressing them for the measurement of global poverty.

Finally, this article contributes to our understanding of the forces shaping the world distribution of income. Global inequalities have undergone profound transformations, including declining poverty (Bourguignon and Morrisson, 2002; Chen and Ravallion, 2010; Hammar and Waldenström, 2020; Page and Pande, 2018), the emergence of a new "global median class" (Lakner and Milanovic, 2016), and rising top income inequality (Chancel and Piketty, 2021). This article studies the role of government transfers in shaping these dynamics. In incorporating public services into poverty statistics, my work also relates to recent efforts at extending poverty measurement to broader dimensions of quality of life beyond market consumption (Alkire, Kanagaratnam, and Suppa, 2021; Baland, Cassan, and Decerf, 2022; World Bank, 2018).

Section 2 presents motivating evidence and the conceptual framework. Section 3 describes

[2] See for instance Paulus, Sutherland, and Tsakloglou (2010), Verbist, Förster, and Vaalavuo (2012), and Wagstaff et al. (2014) on education and health and Aaberge et al. (2010), 2019 on local government services.

the data sources and methodology. Section 4 presents the main results. Section 5 extends the analysis to the study of government productivity. Section 6 turns to the welfare value of public education and healthcare. Section 7 discusses how public services can rationalize discrepancies between surveys and national accounts aggregates. Section 8 concludes.

2. Motivating Evidence and Conceptual Framework

This section presents motivating evidence for studying the distribution of public services (section 2.1) and introduces the general framework used in the paper (section 2.2).

2.1. Motivating Evidence

I start by providing motivating evidence for incorporating the consumption of public services in poverty and inequality statistics. I establish two simple stylized facts.

2.1.1. Public and Private Services Are Substitutes

The standard approach to measuring poverty and inequality focuses on household disposable income or consumption (disposable income minus saving). Disposable income is the sum of labor and capital incomes, minus direct taxes paid, plus cash transfers received. By definition, it excludes public services, which amounts to assuming that their value is exactly zero.

This assumption can lead to implausible conclusions when analyzing the incidence of government redistribution on inequality. Consider a government that decides to fully subsidize healthcare, bringing down private healthcare expenditure to zero. Theoretically, individual incomes should be adjusted by adding the new government in-kind transfer. Yet, poverty and inequality will remain unchanged in standard statistics, because the value of subsidized healthcare is ignored. More generally, every policy subsidizing a service that was privately bought will be measured as having no incidence on poverty or inequality.

Figure 2a provides suggestive evidence that this channel is empirically relevant. There is a strong negative correlation between public health spending and the share of households pushed into extreme poverty by out-of-pocket healthcare expenditure. In Bangladesh, where public health expenditure is 0.5% of net national income, 7% of households see their daily expenditure fall below PPP $3.65 because of private health spending. In South Africa, where public health expenditure is almost 7% of NNI, less than 0.3% of the population ends up poor because of out-of-pocket health spending. Private and public services are therefore not independent. In-kind transfers allow poor households to save money, and not accounting for such money leads to overestimating poverty in countries with large welfare states.

2.1.2. Public Services Matter for Non-Monetary Dimensions of Quality of Life

Public services also contribute to improving non-monetary dimensions of well-being. The need to go beyond strictly monetary measures of poverty has been increasingly recognized in the past decades. Researchers and international organizations have started developing multidimensional poverty indicators, which often involve aggregating measures of well-being across a number of domains. For instance, Alkire, Kanagaratnam, and Suppa (2021) propose a poverty indicator that combines measures of deprivation in health, education, housing, and other basic goods.

Figure 2b provides suggestive evidence that accounting for in-kind transfers contributes to bridging the gap between monetary and multidimensional poverty statistics. The x-axis plots public spending on education, health, and housing and community amenities. The y-axis plots the difference between the multidimensional and monetary poverty rates. There is a strong negative correlation: multidimensional poverty is lower than monetary poverty in countries with large welfare states, while it is higher in countries with low public spending. This suggests that in-kind transfers improve the well-being of poor households in dimensions of quality of life that are not captured by monetary poverty indicators. By incorporating measures of the consumption of public services into income distribution statistics, this article contributes to bridging the gap between these two approaches to poverty measurement. Doing so also contributes to solving discrepancies between survey and national accounts estimates of living standards, as I discuss in section 7.

2.2. Conceptual Framework

I study the consumption of public services by combining data on their value and their distributional incidence. Consider individual i receiving pretax income m_i, paying taxes $\tau(m_i)$, and receiving cash $c(m_i)$ and in-kind transfers $g(m_i)$ from the government. Their posttax income is:

$$\underbrace{y_i}_{\substack{\text{Posttax}\\ \text{Income}}} = \underbrace{m_i}_{\substack{\text{Pretax}\\ \text{Income}}} - \underbrace{\tau(m_i)}_{\text{Taxes}} + \underbrace{c(m_i)}_{\substack{\text{Cash}\\ \text{Transfers}}} + \underbrace{g(m_i)}_{\substack{\text{In-Kind}\\ \text{Transfers}}} \tag{1}$$

The value of in-kind transfers received by i is:

$$g(m_i) = \sum_j g^j(m_i) = \sum_j \underbrace{G^j}_{\text{Value}} \times \underbrace{\gamma^j(m_i)}_{\text{Progressivity}} \tag{2}$$

G^j is the total value of public service j in a given country-year. $\gamma^j(m_i)$ is the share of G^j received by individual i. By definition, $\gamma^j(m_i) \in [0,1]$.

In my benchmark specification, I define G^j as total general government expenditure on j (education or healthcare), in line with the usual approach in the literature (e.g., Auerbach, Kotlikoff, and Koehler, 2023; Auten and Splinter, 2024; Piketty, Saez, and Zucman, 2018). The advantage is consistency with the national accounts, which has been increasingly recognized as a desirable property of income distribution statistics (Statistics Canada, 2019; Stiglitz, Sen, and Fitoussi, 2009). Allocating public spending to individuals thus allows constructing estimates of redistribution that are consistent with global GDP growth.

Yet, valuation at cost of provision suffers from two limitations. First, national accounts may not adequately measure government output in the presence of cross-country and time variations in public sector productivity. Second, individuals' willingness to pay for public services may differ from their cost of provision. In sections 5 and 6, I thus turn to two complementary measures of G^j: one that anchors cost of provision on education and health outcomes, and one that values public services based on returns to schooling and gains in life expectancy enabled by the education and healthcare systems.

3. Data Sources and Methodology

I now turn to the construction of a new database on the provision of public services worldwide. I first cover the distribution of pretax income (section 3.1), followed by the construction of public spending series (section 3.2). I then present the data sources used to estimate the distributional incidence of education, healthcare, and social assistance (sections 3.3 to 3.5). Table 1 provides information on the geographical coverage of the database and summary statistics. Appendix A presents additional methodological details.

3.1. Pretax Income

The starting point is the distribution of pretax income. Data come from the World Inequality Database (WID: see Chancel and Piketty, 2021). The WID compiles estimates from various studies to cover the distribution of pretax income in all countries in the world since 1980. Average income in each country-year matches net national income per capita (NNI: GDP, minus consumption of fixed capital, plus net foreign income). The income concept is pretax national income, that is, income before accounting for the operation of the tax-and-transfer system, but after accounting for the operation of the pension and unemployment systems.

3.2. Public Spending

The first step is to estimate how much governments spend and on which types of policies. Somewhat surprisingly, there exists no single data source covering public expenditure on education, healthcare, and social assistance for all countries in the world. I have thus constructed a new database on general government expenditure for the purpose of this paper.

I exploit data series from two types of sources. A first set of sources provide information on the full breakdown of government expenditure by function. These mainly include series compiled by international providers such as the IMF, Eurostat, the OECD, CEPAL, and the United Nations. A second set of sources cover specific dimensions of expenditure. These include the UNESCO for education, the World Health Organization for healthcare, as well as a number of other sources. In total, I assemble and link series from eighteen sources, allowing me to construct comparable series covering 174 countries since 1980 (167 for social assistance).

3.3. Progressivity of Education

Following the existing literature, I distribute education spending to children attending school in the household (e.g., Bruil et al., 2022; Germain et al., 2021; Lustig, 2018). The data source is a new microdatabase I have constructed for this paper. It consists of about 1,300 nationally representative surveys fielded from 1980 to 2021 in 145 countries, covering 94% of the world's population. These surveys are the ones routinely run by statistical institutes to track the evolution of poverty and inequality, such as the CPS in the United States. They record detailed information on the structure of the household, school attendance, age, and total household income or consumption. The World Bank has assembled a large number of these surveys for most countries around the world since the 1980s. I was able to access this microdatabase, which constitutes my primary data source. To further expand coverage, I complement it with additional country-specific sources covering 51 countries.

I calculate the transfer received by income group p in country c at time t as:

$$g_{pct}^{\text{educ}} = n_{pct}^{pri} g_{pct}^{pri} + n_{pct}^{sec} g_{pct}^{sec} + n_{pct}^{ter} g_{pct}^{ter} \tag{3}$$

Where n_{pct}^{k} denotes the number of children in school at level k, g_{pct}^{k} denotes average spending per child on function k, and $k \in \{pri, sec, ter\}$ refers to primary, secondary, and tertiary education. I thus allocate to each child attending school at level k the per-pupil expenditure on education at this level. Data on the costs of primary, secondary, and tertiary education per child come from the public spending database constructed in this paper. The number of children in school by level and per-capita household income is computed from the microdatabase.

Based on equation 3, it is useful to notice that the distribution of education spending is driven by three factors. The first one is inequality in access to schooling: education is more unequally distributed in countries where only children from high-income households attend school. The second one is spending inequality: countries differ in relative expenditure on primary, secondary, and tertiary education per child. The third one relates to demography: education transfers received are higher in households with more children in school. The first two factors tend to increase inequality in education spending. The third factor often makes education expenditure more progressive, because low-income households tend to have more children.

3.4. Progressivity of Healthcare

I distribute health expenditure proportionally to intensity of use of the public healthcare system. I rely on four sources. The first data source is the Commitment to Equity Institute (CEQ) database, which compiles estimates of tax-and-transfer progressivity from existing studies in 20 countries (see Lustig, 2018). These estimates are typically constructed from survey microdata covering frequency of use of public healthcare, such as the number of visits made to a public health institution in the past month. The second data source consists in 36 additional surveys, which I have manually collected from country-specific data portals and harmonized. The third data source is the World Health Survey, which again provides comparable information on healthcare use intensity for 53 additional countries. Together, these sources allow me to cover 109 countries representative of 86% of the world's population.

As in the case of education, the progressivity of public healthcare expenditure depends on two main factors. The first one is access: it may be more difficult for low-income households to access a public healthcare provider. The second is health inequality: low-income individuals may have greater healthcare needs and hence use the healthcare system more intensively. The first factor increases inequality in health spending, while the second one reduces it.

3.5. Progressivity of Social Assistance

Social assistance expenditure includes expenditure on cash transfers and in-kind social benefits such as food stamps.[3] These transfers are those typically included in traditional measures of government redistribution (e.g., OECD, 2011). I distribute these transfers to their beneficiaries. The data sources are Piketty, Saez, and Zucman (2018) for the United States, Blanchet, Chancel, and Gethin (2022) for European countries, the CEQ database, and the World Bank's ASPIRE database. These sources provide information on the share of social assistance transfers received by income decile or percentile in 130 countries (91% of the world's population).[4]

3.6. Imputation of Missing Data

This paper aims to study how government redistribution shapes global inequality. Yet, while public spending data is almost complete, the coverage of distributional incidence series is incomplete. I thus impute missing data using two simple assumptions (1) for missing countries and (2) for missing years in a given country. Overall, the main results of this paper are largely

[3] I follow the definition of social assistance expenditure adopted by the COFOG classification of the UN System of National Accounts. Social assistance consists in non-contributory transfers to households in need. It excludes pensions, unemployment insurance, and workfare programs.

[4] In each case, I only distribute social assistance expenditure and exclude pensions and unemployment benefits, given that these transfers are already included in pretax income (see section 3.1).

insensitive to alternative imputation assumptions.

For missing countries, I use the average distributional incidence profile observed across all country-years. This imputation has limited implications, because most countries are already covered for at least one year (see Table 1).

For missing years in a given country, I extrapolate series backwards and forwards using the last incidence profile observed, so as to cover the whole 1980-2022 period. This has somewhat stronger implications for healthcare, for which only one year of data is usually available: it amounts to assuming that progressivity has remained constant over time. I discuss the plausibility of this assumption in the next section.

3.7. Validation

3.7.1. Sources of Bias

The objective of this paper is to study the world distribution of income since 1980, yet data quality and coverage are far from perfect. There are four main sources of concern.

Concern 1: Mismeasurement First, one may be concerned with the quality of the data used to measure education and healthcare use intensity. School attendance and visits to healthcare providers are only imperfect measures of the actual use of education and healthcare services. Whether they provide a good first-order approximation remains an open question.

Concern 2: Regional Inequality Public education and healthcare transfers may vary not only depending on intensity of use, but also across subnational regions. Poorer regions may benefit from lower spending, leading education and healthcare progressivity to be overestimated.

Concern 3: Private Education and Healthcare These estimates do not account for the fact that some households rely on private schools and hospitals, who typically benefit from lower government funding than public institutions. Ignoring this channel will lead to *underestimating* progressivity, since private services are disproportionately used by high-income households in virtually all countries with available data (e.g., Lustig, 2018).

Concern 4: Trends in Progressivity My estimates capture trends in progressivity imperfectly. The microdatabase does cover trends in schooling inequality, but not for all countries and generally not for the whole time period. Data on the distributional incidence of healthcare are even scarcer, typically covering one year per country (see Appendix Table B8).

3.7.2. Validation of Methodology

In the face of these challenges, I conduct three separate validation exercises.

Validation 1: Comparison with High-Quality Studies First, I compare my estimates with those of three high-quality studies covering France (Germain et al., 2021), the Netherlands (Bruil et al., 2022), and South Africa (Gethin, 2024), which use detailed administrative data to measure the distributional incidence of government transfers. In particular, they use much more precise measures of the consumption of public services (concern 1), incorporate regional spending inequality (concern 2), and account for private education and healthcare (concern 3).[5] Figure 3 compares government transfers received by the bottom 50% in these detailed studies and in this paper. My estimates closely align with those of these studies.

Validation 2: Comparison with CEQ Estimates Second, the CEQ database also provides high-quality estimates of the distribution of education and healthcare that account for private consumption and geographical spending inequalities (concerns 2 and 3). Again, my simplified estimates closely align with those of the CEQ (Figures 4a and 4b). If anything, I slightly underestimate redistribution in the average country.

[5] Germain et al. (2021) and Bruil et al. (2022) exploit comprehensive administrative microdata on school attendance and actual health expenditure. Gethin (2024) relies on household surveys providing rich information on school attendance and use of different types of healthcare services, which are combined with administrative budget data on education and healthcare expenditure by function and province.

Validation 3: Investigating Trends in Progressivity Third, one might be concerned about unobserved trends in progressivity. Put simply, I assume that $\gamma^j(m_i)$ has remained constant when no longitudinal data is available. If progressivity has declined, then I will be overestimating improvements in public services received by the global poor.

To make progress in better understanding trends in progressivity, I compare the first and last years available in countries with historical data. For education, the microdatabase constructed in this paper covers time horizons exceeding five years in 128 countries. For healthcare, I collect and harmonize additional historical surveys covering eight large countries, which I complement with CEQ data covering six additional countries. There has been little change in progressivity in the majority of countries (Figures 4c and 4d). If anything, the data suggest some improvement: 84 out of 128 countries saw an increase in education progressivity, while 8 out of 13 countries saw an increase in healthcare progressivity.

A last piece of interesting evidence comes from the Global Burden of Disease study (GBD, 2022), which provides longitudinal subnational data on healthcare quality in seven countries (see section 5 for a longer discussion). In every single of these countries, there has been a decline in regional inequality in the quality of healthcare since the 1990s (Appendix Figure A1).

Together, these results demonstrate that the simplified estimates of redistribution developed in this paper provide an excellent first-order approximation of cross-country and time variations in the distributional incidence of government transfers around the world.

4. Main Results

This section presents the main results, focusing on the valuation of public services at cost of provision. Section 4.1 studies cross-national variations in government redistribution. Section 4.2 quantifies the incidence of public services on the world distribution of income.

4.1. Government Redistribution Around the World

4.1.1. Cross-Country Differences in Public Spending

I start by documenting levels and trends in public spending. Figure 5 plots the evolution of government expenditure on social assistance, education, and healthcare by country income group from 1980 to 2022. In 2022, expenditure on social assistance, education, and healthcare was about 7% of national income in low-income countries, 8-12% in middle-income countries, and 23% in high-income countries. Low-income countries also dedicate a lower fraction of expenditure to social assistance, which represents about 15% of total transfers compared to a third in high-income countries. Ignoring the consumption of public services would thus lead to miss a substantial fraction of redistribution, especially in poorer countries.

Government redistribution has increased. Total spending grew from 5.2% to 7% of NNI in low-income countries—a 35% increase. It expanded at about the same pace in middle-income and high-income countries. The increase in public spending is visible across all three functions of government, although it was primarily driven by education in low-income countries and social assistance and healthcare in high-income countries.

4.1.2. Cross-Country Differences in Progressivity

I now turn to variations in the distributional incidence of cash and in-kind transfers. Figure 6 plots the average share of pretax income and government transfers received by the bottom 50%, separately by country income group.

Government transfers unambiguously reduce inequality. The average share of pretax income received by the poorest 50% varies from 14% in low-income countries to 16% in high-income countries. Meanwhile, 50-60% of transfers are received by the bottom 50%. In other words, government redistribution is well approximated by a lump sum allocation in low-income countries and is even more progressive than a lump sum in high-income countries.

Social assistance is the most progressive function of government, which is to be expected given the often explicitly pro-poor design of the corresponding programs (such as conditional cash transfers or food stamps). In low-income countries, however, social assistance transfers are not particularly well targeted at low-income households.

Education is less progressive than cash transfers but still substantially reduces inequality: 45% of spending accrues to the bottom 50% in low-income countries, compared to 55% in high-income countries. While such degree of progressivity might seem surprising, it is useful to remind the reader that it is the product of countervailing forces (see section 3.3). On the one hand, children from high-income households tend to stay longer in school. On the other hand, low-income households tend to have more children. These effects more or less compensate each other, yielding a quasi-egalitarian distribution of education spending in the average country.

Public healthcare is about as progressive as education, which reflects again the joint effect of two forces. On the one hand, low-income households often have lower access to public healthcare because of user fees or travel distances. On the other hand, they tend to have greater healthcare needs. This second effect dominates in high-income countries, where the distribution of healthcare is more progressive than a lump sum allocation—consistently with evidence on France and the Netherlands using administrative data (Bruil et al., 2022; Germain et al., 2021).

4.1.3. Cross-Country Differences in Redistribution

Combining data on the value and progressivity of cash and in-kind transfers, I obtain estimates of overall redistribution in each country. Figure 7a maps social assistance, education, and healthcare transfers received by the bottom 50% in 2022 as a share of NNI. Redistribution is highest in North America and Western Europe, exceeding 8% of NNI in many countries. It is also high in Latin America and Southern Africa. Government transfers are lower in Asia, falling below 4% in most countries, and lowest in Western, Central, and Eastern Africa.

Figure 7b zooms into transfers received by the bottom 50% in fifteen large countries or regions, which together represent about two-thirds of the world's population. Transfers amount to about 2% of NNI in Pakistan, Bangladesh, and Ethiopia, compared to 15-16% in Western Europe and the United States. The bulk of transfers is made in the form of education and healthcare in low- and middle-income countries. Ignoring public services would thus lead to massively underestimating redistribution in the developing world.

Finally, Figure 8 plots the evolution of redistribution around the world by comparing the share of national income received by the bottom 50% in 1980 (x-axis) and 2022 (y-axis). Redistribution has increased in most countries—122 countries are located above the 45-degree line. This increase has been substantial in several big countries such as China, India, Brazil, and South Africa. Other countries such as Nigeria, Mexico, or Russia saw redistribution stagnate or slightly decrease. Redistribution has strongly declined only in a handful of countries, most of which have undergone prolonged periods of instability, such as Zimbabwe, Sudan, and Haiti.

4.2. Government Redistribution and the World Distribution of Income

4.2.1. Government Redistribution and Global Poverty Reduction

I now turn to the role of government transfers in shaping the world distribution of income. I first convert all pretax incomes and transfers to real 2021 PPP US dollars in each country-year. I then rank all individuals in the world from the poorest to the richest, and aggregate average incomes and transfers received by global income percentile.

Figure 9 plots the evolution of total government transfers received by the world's poorest 20% from 1980 to 2022. Transfers received by the global poor have doubled as share of global GDP. This progress was largely driven by public services, which represent over two-thirds of transfers in 2022. At the same time, global income per capita approximately doubled. The real transfer received by the global bottom 20% was thus multiplied by four (see Appendix Figure A3).

Figure 1 plots the evolution of the real average income of the world's poorest 20%, expressed in

2021 PPP US dollars, before and after adding cash and in-kind transfers to individual incomes. Average pretax income increased by about 80%. This figure rises to about 125% after adding social assistance transfers, and 175% after also accounting for the consumption of public services. By this measure, total government redistribution accounts for about 50% of global bottom 20% income growth. Public education and healthcare account for 30%. This general result is robust to a large set of alternative specifications. Depending on the exact assumptions, data sources, and indicators used, redistribution is found to account for 40-60% of global bottom 20% growth from 1980 to 2022.[6]

4.2.2. Government Redistribution and Global Income Inequality

I now extend the analysis to the overall distribution of global economic growth. Figure 10 plots the real income growth rate experienced by each global income percentile from 1980 to 2022, before and after adding government transfers. Pretax growth have been greatest at the middle and top end of the distribution, generating what has been referred to as the "elephant curve" of global inequality and growth (Lakner and Milanovic, 2016). This pattern reflects the conjunction of trends in inequality between and within countries, including the rise of China and India (middle of the distribution), sluggish economic growth in low-income countries (lower end) and high-income countries (upper-middle section), and skyrocketing top income inequality in many parts of the world (top end).

While this pattern has been extensively documented, little is known of the role played by government redistribution in shaping it. As shown in Figure 10, adding government transfers substantially raises growth rates for all groups within the bottom 60% of the world distribution of income. Education and healthcare alone account for over half of this effect, pointing to the

[6]In particular, Appendix Figure A4 reproduces the analysis using poverty headcount ratios at $2.15 per day in 2017 PPP USD. Appendix Table A1 reports results with a more comprehensive measure of posttax income, drawing on estimates of the distributional incidence of taxes from a companion paper (Fisher-Post and Gethin, 2023), as well as results excluding China, India, or both countries from the analysis. Appendix Table A2 reproduces the analysis using World Bank income distribution data instead of the WID.

fundamental role played by public services in making global economic growth more inclusive.[7]

4.2.3. Government Redistribution and the Geography of Global Inequality

Finally, this new database allows studying how public services shape the geography of world inequality. The upper panel of Table 2 provides a Theil decomposition of global inequality into its between-country and within-country components. Government redistribution unambiguously reduces global inequality, but it also increases the share of global inequality explained by inequalities between countries. The between-country component accounts for 30% of pretax global inequality, compared to 33% after social assistance transfers, and 38% after all transfers. Government redistribution, especially through public services, thus contributes to increasing the weight of national differences in GDP per capita in explaining global inequality.

The lower panel of Table 2 breaks down the geographical location of the world's poorest 20% by world region. Accounting for government redistribution raises the share of the global poor living in India, Pakistan, Bangladesh, Ethiopia, Nigeria, and other Sub-Saharan African countries. It improves the relative positions of low-income households in China, Latin America, and the Western world. These differences are large. For instance, moving from pretax to posttax income increases the share of poor Pakistani households from 17% to 27%—a 50% increase. In Europe and North America, on the contrary, this share drops from 8% to zero.

5. Accounting for Public Sector Productivity

The previous analysis valued public education and healthcare at cost of provision. A natural concern is that public spending may not adequately reflect the output of the public sector, given potentially large variations in government efficiency. In this section, I investigate the

[7] Appendix Figure A5 extends this analysis to global inequality, measured as the ratio of the world's richest 10% to poorest 50% average incomes. Redistribution is found to account for about a third of global inequality reduction since 1980. Appendix Figures A6 and A7 reproduce this analysis using Gini and Theil indices.

implications of accounting for public sector productivity for my analysis.

5.1. Conceptual Framework

Public spending may not accurately reflect levels and trends in the output of the public sector. Put simply, national accounts may be wrong. This concern has led to the emergence of a growing literature seeking to adjust national accounts aggregates to better reflect the quality of education and healthcare (see, e.g., Cutler et al., 2022 for the U.S. healthcare sector).

In order to make progress on this issue, I propose a minimal adjustment to public spending figures that anchors the value of public services to actual education and health outcomes. Formally, let us reexpress the value of public service j as:

$$\underbrace{G^j}_{\text{Value}} = \underbrace{E^j}_{\text{Cost of Provision}} \times \underbrace{\Theta^j}_{\text{Productivity}} \tag{4}$$

Where E^j is total general government expenditure on education or healthcare, and Θ^j is a productivity parameter capturing the fact that for a given cost of provision, individuals may receive services of different quality. $\Theta^j = 1$ is the main specification used until now, while $\Theta^j = 0$ corresponds to a case in which public spending is completely useless.

5.2. Data Sources and Methodology

I propose to estimate Θ^j by benchmarking the productivity of governments around the world to one another (e.g., Adam, Delis, and Kammas, 2011; Herrera and Ouedraogo, 2018). If a government delivers public services of better quality than any other at a given cost, then its productivity is set to $\Theta^j = 1$. All governments with a comparable cost but lower outcomes are then attributed a Θ^j between 0 and 1, based on their distance to this "efficient frontier." The advantage of this approach is its simplicity and transparency: governments delivering better outcomes are considered to be more productive.

For education, I measure the output of the public sector as expected human capital from age 5 onward, estimated by combining data on school attendance and test scores from international databases. For healthcare, the outcome of interest is the healthcare access and quality index provided by the global burden of disease study (GBD, 2022), which ranks healthcare systems from 0 to 100 based on death rates from 32 causes of death that could be avoided by timely and effective medical care. I choose these indicators for two reasons. First, they are among the only ones for which historical data is available for almost all countries in the world. Second, they are direct measures of education and healthcare quality, in contrast to other outcomes such as life expectancy, which is arguably more contaminated by unobserved factors (although these two measures are not immune to this issue either). See Appendix A.3 for more details.

Appendix Figures A8 and A9 plot the evolution of these two indicators by country income group. Enormous progress has been made in the provision of quality education and healthcare. Expected human capital has expanded by about 30% in low-income countries, 50-80% in middle-income countries, and 50% in high-income countries. The quality of healthcare has improved by about 60% in low- and middle-income countries, and 20% in high-income countries.

Estimating government productivity amounts to anchoring cost of provision on these educational and health outcomes. Figure 11a provides a concrete illustration in the case of education. Education spending per child is strongly correlated with expected human capital, but there is also significant variation in educational outcomes for a given level of spending. Country-years that perform best at a given cost are attributed $\Theta^j = 1$: they are at the "efficient frontier." Country-years below the frontier are attributed values of Θ^j ranging from 0 to 1. The trajectory of India is highlighted in red. India lies significantly below the efficient frontier, but this gap has remained relatively constant, implying a Θ^j around 0.6 during this period.

Figure 11b reproduces the same figure for healthcare. There is again a very strong correlation between health expenditure and quality of healthcare, but significant variations at any given level of spending. India falls significantly below the frontier, with Θ^j estimated to range from

0.35 to 0.55 during this period.

5.3. Discussion and Validation

I view these estimates as providing a *lower bound* on government productivity, especially in poor countries, for three main reasons. First, PPP conversion factors already make an adjustment for public sector productivity, so this approach holds the risk of "double-counting" inefficiencies (World Bank, 2013). Second, this methodology implies always reducing transfers in all countries that are not at the frontier ($\Theta \leq 1$). This is equivalent to assuming that governments are never more efficient than the private sector. Third, omitted variable bias implies that productivity is likely to be underestimated in low-income countries, whose lower outcomes are arguably the product of other factors than government performance (such as lower income itself).

A useful way of validating my measures of government productivity is to compare them to existing indicators. Appendix Table A4 shows that education and healthcare productivity are positively correlated with a number of other indicators of government efficiency. This is especially true of healthcare productivity, which is positively associated with the World Bank's composite index of government effectiveness ($\rho = 0.66$), lower corruption ($\rho = 0.58$), and more transparent policy-making ($\rho = 0.43$).

5.4. Main Results

I now turn to the implications of accounting for public sector productivity for the results of this paper. There are two main findings.

First, accounting for productivity magnifies cross-country differences in redistribution. In 2022, average productivity in education and healthcare were 0.63 and 0.59 in low-income countries, compared to 0.72 and 0.93 in high-income countries.[8] As a result, accounting for productivity reduces in-kind transfers by about 40% in low-income countries, compared to 15%

[8]See Appendix Table A3.

in high-income countries.[9] Low-income countries thus appear to suffer from a "triple curse" of redistribution: not only do they spend less on public services and distribute them more unequally, they also provide them less efficiently.

Second, accounting for productivity mechanically reduces in-kind transfers received by the global poor. However, it does not significantly alter the trend; as a result, my results on the role of redistribution in reducing global poverty and inequality remain essentially unchanged.[10]

6. The Welfare Value of Public Education and Healthcare

The previous sections relied on valuing public education and healthcare in a national accounts framework. An open question is whether these measures reflect individuals' willingness to pay for these services. In this section, I make progress in estimating the welfare value of public services. Given the important challenges surrounding this exercise, these estimates should be interpreted with more caution. The main objective is to investigate the sensitivity of my results to this alternative valuation method. Importantly, this approach relies on different methodological principles and data sources than the ones used in my benchmark specification, which makes it an appealing complementary exercise.

6.1. Methodology

6.1.1. Welfare Value of Education

Framework I value public education transfers based on the net present value of expected returns from receiving an additional year of schooling. I make the following simplifying assumptions. There are four levels of education in the economy: no schooling (0), primary

[9] See Appendix Table A5.

[10] Appendix Figure A10 plots the total transfer received by the world's poorest 20% before and after adjusting for productivity. Appendix Figures A11 to A12 turn to the global bottom 20% average income and the distribution of economic growth.

(1), secondary (2), and tertiary education (3). Individuals receive an average income y_{ctps} at time t, depending on their country c, income group p, and level of education s. They can go to school for one year, delivering a return γ_{cs} when they start working. They receive these monetary benefits of education every year for 45 years of working life (say, from age 20 to 65), discounting them at a rate δ. The economy grows at a constant rate g. With these assumptions, I show in Appendix A.4 that the net present value of the public education system is:

$$g_{ctp}^{\text{educ}} = \sum_{s=1}^{3} n_{ctps} \left(\sum_{i=t}^{t+45} \frac{(1+g)^i \gamma_{cs} y_{ctp,s-1}}{(1+\delta)^i} \right) \tag{5}$$

With n_{ctps} the number of children belonging to income group p that are enrolled in public schools at level s. Intuitively, the real value of public education is higher for low-income households when many of their children are in school (n_{ctps}), the return to schooling is large (γ_{cs}), expected income absent additional schooling is also large ($y_{ctp,s-1}$), the economy is growing fast (g), and future benefits from education are discounted at a lower rate (δ).

Data Sources The number of students in school by level and income group come from the microdatabase collected in this paper. Returns to schooling γ_{cs} come from a companion paper (Gethin, forthcoming). Data on expected incomes by education level and income group $y_{pt,s-1}$ are constructed by combining data on average incomes by education level from Gethin (forthcoming), intergenerational mobility curves from van der Weide et al. (2024), and observed and projected growth rates over 1980-2050 from the OECD. The discount rate is set at $\delta = 5\%$.

Discussion I view this measure as providing a lower bound for two reasons. First, this approach focuses exclusively on the monetary benefits of education, while schooling also has large positive effects on other dimensions of quality of life such as mental health, crime, or fertility (Oreopoulos and Salvanes, 2011). Second, education not only increases individual incomes, but also generates positive human capital externalities (e.g., Gennaioli et al., 2013).

6.1.2. Welfare Value of Healthcare

Framework In the same spirit, I estimate the welfare value of public healthcare based on the monetary value of additional years of life expectancy enabled by the healthcare system. Let V_{ctp} be the value of a year of life for percentile p in country c at time t. Individuals can expect to live T_{ct} years longer, L_{ct} years of which can be attributable to public healthcare. An average individual thus receives every year until their death an annual life expectancy gain of $\frac{L_{ct}}{T_{ct}}$. The net present value of the public healthcare system is:

$$g_{ctp}^{\text{heal}} = \zeta^{publ} \sum_{i=0}^{T_{ct}} \frac{\frac{V_{cip} L_{ct}}{T_{ct}}}{(1+\delta)^i} \tag{6}$$

Assuming that the public healthcare system accounts for a fraction ζ^{publ} of life expectancy gains.

One would like to estimate V_{ctp}, the value of a statistical life year, for every country and income group in the sample. Unfortunately, such estimates are not available. However, a lower bound on this value is the expected income that individuals can hope to receive (Hammitt and Robinson, 2011). This is a lower bound since individuals receive utility not only from income but also from leisure. This additional income is only realized in $T_{ct} - L_{ct}$, when the individual starts benefiting from greater life expectancy, and is received until they die, in T_{ct}. The net present value of the healthcare system can then be reexpressed as:

$$g_{ctp}^{\text{heal}} = \frac{1}{T_{ct}} \zeta^{publ} \sum_{i=T_{ct}-L_{ct}}^{T_{ct}} \frac{y_{cip}}{(1+\delta)^i} \tag{7}$$

With y_{cip} the average income that income group p in country c can expect to receive at time i. In summary, the value of the healthcare system is the discounted sum of additional cash flows that individuals receive from living longer thanks to public healthcare.

Data Sources Data on remaining years of life is calculated as the difference between life expectancy and the average age, as in common in the literature. ζ^{publ} is proxied by the share of public spending in total healthcare expenditure. Data on average income by income group y_{pci} come from this paper. The discount rate is set at $\delta = 5\%$, as for education. Finally, life expectancy gains enabled by the public healthcare system L_{ct} are estimated by regressing the Healthcare Access and Quality index on life expectancy and normalizing it using the coefficient obtained. The resulting index ranges from less than 1 (Eritrea) to 34 (Finland): an average Finn would live 34 years less if there was no healthcare system at all. I provide a longer discussion and validation of this indicator in Appendix A.3.3. In particular, I find that it is closely tracks available estimates of medical progress for the United States (Cutler et al., 2022).

6.2. Main Results

Four main results emerge from the analysis of these additional indicators.

First, I compare cost-of-provision and welfare estimates of public education and healthcare in each country in Figure 12.[11] The two measures are positively correlated ($\hat{\rho} = 0.3$ for education and 0.6 for healthcare). This suggests that cost of provision provides an imperfect, but still useful first-order approximation of cross-country differences in the value of public services.

Second, the welfare value of public services is higher than its cost of provision in almost all countries. There are large variations in this gap, however. The net present value of education is low in India and South Africa, two countries that are well-known for having low quality of education (Mlachila and Moeletsi, 2019; Muralidharan and Sundararaman, 2015). The United States is one of the countries closest to the 45-degree line on healthcare, consistently with evidence on the low welfare value of Medicaid (Finkelstein, Hendren, and Luttmer, 2019).

Third, the welfare value of education and healthcare has considerably increased, especially in low- and middle-income countries. In low-income countries, it has been multiplied by four as a share of national income (see Appendix Figures A13 and A14).

Fourth, I extend this analysis to the study of the world distribution of income. In my benchmark

[11] Appendix Figure A15 reproduces this figure for the transfer received by the bottom 50%.

specification, the welfare value of public services received by the world's poorest 20% is about two times higher than their cost of provision and has risen at a slightly faster pace. As a result, government redistribution is found to account for 65% of growth among the world's poorest 20% when public services are valued at their net present value (45%-65% depending on the exact data sources and specification).[12]

7. Surveys versus GDP: Public Services and Measurement Discrepancies in Global Poverty Statistics

I conclude this paper with a discussion of how the consumption of public services can elucidate puzzling gaps between survey and national accounts aggregates documented in the literature.

7.1. Motivation

A major debate in development economics centers around whether national accounts or surveys should be used in priority to measure economic development. For reasons that are not well understood, persistent discrepancies between GDP and survey incomes can lead to conflicting conclusions on the evolution of living standards in the past decades (Deaton, 2005).

Recent studies point to the superiority of GDP. Combining data from various sources, Pinkovskiy and Sala-i-Martin (2016) provide evidence that GDP correlates much more significantly with satellite-recorded nighttime lights than survey means. It also accounts for a greater fraction of variations in several indicators of quality of life, such as life expectancy, access to safe water, and primary school enrollment. Most importantly, the difference between GDP and survey means is positively associated with achievements on these indicators. In other words, "countries with higher and growing well-being tend to suffer from progressively greater mismeasurement of

[12]Appendix Figure A16 plots total government transfers received by the world's poorest 20%. It also reports estimates of a conservative lower bound on the welfare value of public services in which g is set to zero. Appendix Tables A6 and A7 extend this analysis to global bottom 20% growth and the global poverty headcount ratio.

income by surveys." While the authors suggest that this finding could be due to the complexity of survey questionnaires, the exact reasons underlying this result remain unclear.

There is one natural candidate for explaining this discrepancy: public services. As previously mentioned, surveys entirely miss government-provided education, health, and other services, which are not bought on a market and are thus absent from standard consumption measures. Arguably, these services play a key role in improving quality of life in the exact dimensions studied by Pinkovskiy and Sala-i-Martin (2016). The share of national income spent on public services also appears to have risen in the past decades, which could explain why surveys and GDP have become increasingly disconnected from each other.

7.2. Public Services and the GDP-Survey Gap

I investigate this possibility in Table 3. As in Pinkovskiy and Sala-i-Martin (2016), I regress four quality of life indicators on the gap between GDP and survey means: youth literacy, secondary school enrollment, infant mortality, and life expectancy. I then compare the coefficient obtained before and after controlling for public spending on education and healthcare.

In line with Pinkovskiy and Sala-i-Martin (2016), I find that the gap between GDP and surveys tends to be positively correlated with greater quality of life, both before and after adding country fixed effects (panels A and B). However, controlling for spending on education and healthcare considerably reduces the size of the coefficient, which becomes statistically non-significant or even negative in most specifications. In other words, the reason why educational and health outcomes are positively associated with the GDP-survey gap is that countries with such growing gaps have been investing more in public education and healthcare. Hence, the reason why GDP estimates track indicators of quality of life better than surveys is that they include the consumption of public services while surveys do not. In directly incorporating this "missing consumption" into poverty and inequality statistics, I attempted in this paper to correct some of the discrepancies between these two approaches to the measurement of living standards.

8. Conclusion

This article represented an attempt at incorporating measures of public service delivery in global poverty statistics, focusing on education and healthcare. I showed that doing so leads to a more positive view of poverty reduction since 1980, because public services are strongly progressive and governments have been increasingly investing in them. Nonetheless, the share of the world's GDP accruing to the global poor remains limited, because low-income countries provide public services in lower quantities, less progressively, and less efficiently than in the rich world. There is space for improvement in all three of these dimensions of redistribution. Enhancing tax revenue, improving equity in access to public services, and raising government productivity should be seen as complementary levers in the fight against global poverty.

This article has taken a broad historical perspective on poverty reduction, yet much remains to be done to better track public goods delivery. First, there is an urgent need for more transparency on what governments actually do. The data exploited in this article cover spending on large categories, such as education, healthcare, and social assistance, with only limited information on the underlying policies. Unfortunately, information on these policies remains scarce; even when it exists, it often ends up buried under a multitude of documents published by different institutions. The publication of regular reports consolidating and harmonizing data on government budgets, with precise information on the corresponding policies, should be viewed as a priority not only for government accountability, but also for the measurement of poverty and inequality. Too often, researchers and statistical institutes face no other option than to ignore public services, simply because of a lack of data on what these services actually are.

Second, more attention should be given to public services in the design of living standards surveys. Surveys routinely fielded by statistical institutes spend considerable time and effort compiling data on household expenditure, yet the information that they collect on access to basic public services remains rudimentary at best. Adding regular questions on both objective indicators and subjective perceptions of public service delivery would allow for a more complete view of the well-being of low-income households. These questions should be designed in ways that would make them directly comparable with spending data on the different kinds of public services provided by governments.

Third, much more research should be conducted on how individuals actually value public services, not only in comparison to cash but also in comparison to one another. Under which conditions do households prefer to receive a transfer in the form of public healthcare, rather than education or cash? How do these priorities vary across countries, over time, and throughout the income distribution? Evidence on these questions remains scarce. Designing surveys eliciting such preferences would represent an important contribution to our understanding of the role of public services in reducing global poverty. Ideally, specific modules could be directly added to the questionnaires of living standards surveys, so as to regularly collect information on citizens' needs and priorities when it comes to public goods delivery.

References

Aaberge, Rolf, Manudeep Bhuller, Audun Langørgen, and Magne Mogstad (2010). "The distributional impact of public services when needs differ". In: *Journal of Public Economics* 94.1, pp. 9–10.

Aaberge, Rolf, Lasse Eika, Audun Langørgen, and Magne Mogstad (2019). "Local governments, in-kind transfers, and economic inequality". In: *Journal of Public Economics* 180.1, pp. 1–22.

Adam, Antonis, Manthos Delis, and Pantelis Kammas (2011). "Public sector efficiency: leveling the playing field between OECD countries". In: *Public Choice* 146.1, pp. 163–183.

Alkire, Sabina, Usha Kanagaratnam, and Nicolai Suppa (2021). "The Global Multidimensional Poverty Index (MPI) 2021". Oxford Poverty and Human Development Initiative.

Altinok, Nadir, Noam Angrist, and Harry A. Patrinos (2018). "Global Data Set on Education Quality (1965-2015)". World Bank Policy Research Working Paper No. 8314.

Asher, Sam, Tobias Lunt, Ryu Matsuura, and Paul Novosad (2021). "Development Research at High Geographic Resolution: An Analysis of Night Lights, Firms, and Poverty in India using the SHRUG Open Data Platform". In: *World Bank Economic Review* 35.4, pp. 845–871.

Auerbach, Alan J., Laurence J. Kotlikoff, and Darryl Koehler (2023). "US Inequality and Fiscal Progressivity: An Intragenerational Accounting". In: *Journal of Political Economy* 131.5, 1249–1293.

Auten, Gerald and David Splinter (2024). “Income Inequality in the United States: Using Tax Data to Measure Long-Term Trends”. In: *Journal of Political Economy* 132.7, 2179–2227.
Bachas, Pierre, Matthew H. Fisher-Post, Anders Jensen, and Gabriel Zucman (2022). “Globalization and Factor Income Taxation”. NBER Working Paper 29819.
Bachas, Pierre, Lucie Gadenne, and Anders Jensen (2024). “Informality, Consumption Taxes and Redistribution”. In: *Review of Economic Studies* 91.5, 2604–2634.
Baland, Jean-Marie, Guilhem Cassan, and Benoit Decerf (2022). “Integrating Mortality into Poverty Measurement through the Poverty Adjusted Life Expectancy Index”. World Bank Policy Research Working Paper.
Bharti, Nitin and Li Yang (2023). “Human Capital Accumulation in China and India in 20th century”. Working Paper.
Blades, Derek (2013). “Health and Education”. In: *Measuring the Real Size of the World Economy: The Framework, Methodology, and Results of the International Comparison Program – ICP*. Ed. by World Bank. The World Bank. Chap. 11, pp. 301–318.
Blanchet, Thomas, Lucas Chancel, Ignacio Flores, and Marc Morgan (2021). *Distributional National Accounts Guidelines: Methods and Concepts Used in the World Inequality Database*. Tech. rep.
Blanchet, Thomas, Lucas Chancel, and Amory Gethin (2022). “Why Is Europe More Equal Than the United States?” In: *American Economic Journal: Applied Economics* 14.4, pp. 1–40.
Bourguignon, François and Christian Morrisson (2002). “Inequality among World Citizens: 1820-1992”. In: *American Economic Review* 92.4, pp. 727–744.
Bruil, Arjan, Céline van Essen, Wouter Leenders, Arjan Lejour, Jan Mölhmann, and Simon Rabaté (2022). “Inequality and Redistribution in the Netherlands”. Working Paper.
Chancel, Lucas, Denis Cogneau, Alix Myczkowski, and Anne-Sophie Robilliard (2023). “Income Inequality in Africa, 1990-2019: Measurement, Patterns, Determinants”. In: *World Development* 163.1, pp. 106–162.
Chancel, Lucas and Thomas Piketty (2019). “Indian Income Inequality, 1922-2015: From British Raj to Billionaire Raj?” In: *Review of Income and Wealth* 65.S1, S33–S62.
– (2021). “Global Income Inequality, 1820–2020: the Persistence and Mutation of Extreme Inequality”. In: *Journal of the European Economic Association* 19.6, 3025–3062.
Chen, Shaohua and Martin Ravallion (2010). “The Developing World Is Poorer Than We Thought, But No Less Successful In the Fight Against Poverty”. In: *Quarterly Journal of Economics* 125.4, pp. 1577–1625.
Chong, Alberto, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer (2014). “Letter Grading Government Efficiency”. In: *Journal of the European Economic Association* 12.2, pp. 277–299.
Clay, Karen, Jeff Lingwall, and Melvin Jr. Stephens (2021). “Laws, educational outcomes, and returns to schooling evidence from the first wave of U.S. state compulsory attendance laws”. In: *Labour Economics* 68.1, pp. 1019–1035.
Cutler, David M., Kaushik Ghosh, Kassandra L. Messer, Trivellore Raghunathan, Allison B. Rosen, and Susan T. Stewart (2022). “A Satellite Account for Health in the United States”. In: *American Economic Review* 112.2, pp. 494–533.
De Rosa, Mauricio, Ignacio Flores, and Marc Morgan (2022). “More Unequal or Not as Rich? Revisiting the Latin American Exception”. Stone Center for Inequality Studies Working Paper Series.

Deaton, Angus (2005). “Measuring Poverty in a Growing World (Or Measuring Growth in a Poor World)”. In: *The Review of Economics and Statistics* 87.1, pp. 1–19.
Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer (2019). “The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment”. In: *Journal of Political Economy* 127.6, pp. 2836–2874.
Fisher-Post, Matthew and Amory Gethin (2023). “Government Redistribution and Development: Global Estimates of Tax-and-Transfer Progressivity, 1980-2019”. Working Paper.
GBD (2022). “Assessing performance of the Healthcare Access and Quality Index, overall and by select age groups, for 204 countries and territories, 1990–2019: a systematic analysis from the Global Burden of Disease Study 2019”. In: *Lancet Glob Health* 10.12, pp. 1715–1743.
Gennaioli, Nicola, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer (2013). “Human Capital and Regional Development”. In: *The Quarterly Journal of Economics* 128.1, 105–164.
Germain, Jean-Marc, Lucas Chancel, Mathias André, and Thomas Blanchet (2021). *Rapport du groupe d’experts sur la mesure des inégalités et de la redistribution*. Tech. rep. INSEE.
Gethin, Amory (2024). “Who Benefits from Public Services? Novel Evidence and Implications for Inequality Measurement”. Working Paper.
– (forthcoming). “Distributional Growth Accounting: Education and the Reduction of Global Poverty, 1980-2022”. Quarterly Journal of Economics.
Hammar, Olle and Daniel Waldenström (2020). “Global Earnings Inequality, 1970-2018”. In: *The Economic Journal* 130.632, pp. 2526–2545.
Hammitt, James K. and Lisa A. Robinson (2011). “The Income Elasticity of the Value per Statistical Life: Transferring Estimates between High and Low Income Populations”. In: *Journal of Benefit-Cost Analysis* 2.1, pp. 1–27.
Hanushek, Eric A., Jens Ruhose, and Ludger Woessman (2017). “Knowledge Capital and Aggregate Income Differences: Development Accounting for US States”. In: *American Economic Journal: Macroeconomics* 9.4, pp. 184–224.
Hendren, Nathaniel and Ben Sprung-Keyser (2020). “A Unified Welfare Analysis of Government Policies”. In: *The Quarterly Journal of Economics* 135.3, pp. 1209–1318.
Herrera, Santiago and Abdoulaye Ouedraogo (2018). “Efficiency of Public Spending in Education, Health, and Infrastructure”. World Bank Policy Research Paper.
Heston, Alan (2013). “Government Services: Productivity Adjustments”. In: *Measuring the Real Size of the World Economy: The Framework, Methodology, and Results of the International Comparison Program – ICP*. Ed. by World Bank. The World Bank. Chap. 11, pp. 301–318.
Jensen, Anders (2022). “Employment Structure and the Rise of the Modern Tax System”. In: *American Economic Review* 112.1, pp. 213–234.
Lakner, Christoph and Branko Milanovic (2016). “Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession”. In: *The World Bank Economic Review* 30.2, pp. 203–232.
Lustig, Nora, ed. (2018). *Commitment to Equity Handbook: Estimating the Impact of Fiscal Policy on Inequality and Poverty*. Brookings Institution Press.
Mauro, Paolo, Rafael Romeu, Ariel Binder, and Asad Zaman (2015). “A modern history of fiscal prudence and profligacy”. In: *Journal of Monetary Economics* 76.1, pp. 55–70.

Mlachila, Montfort and Tlhalefang Moeletsi (2019). “Struggling to Make the Grade: A Review of the Causes and Consequences of the Weak Outcomes of South Africa’s Education System”. IMF Working Paper.

Montenegro, Claudio E. and Harry Anthony Patrinos (2021). “A data set of comparable estimates of the private rate of return to schooling in the world, 1970–2014”. In: *International Journal of Manpower* 1.1, pp. 1–20.

Muralidharan, Karthik and Venkatesh Sundararaman (2015). “The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India”. In: *The Quarterly Journal of Economics* 130.3, 1011–1066.

Murray, C. J. L., R. Govindaraj, and P. Musgrove (1994). “National health expenditures: a global analysis”. In: *Bull World Health Organ* 72.4, pp. 623–637.

OECD (2011). *Divided we stand: Why inequality keeps rising?* OECD Publishing.

Oreopoulos, Philip and Kjell G. Salvanes (2011). “Priceless: The Nonpecuniary Benefits of Schooling”. In: *The Journal of Economic Perspectives* 25.1, pp. 159–184.

Page, Lucy and Rohini Pande (2018). “Ending Global Poverty: Why Money Isn’t Enough”. In: *Journal of Economic Perspectives* 32.4, pp. 173–200.

Paulus, Alari, Holly Sutherland, and Panos Tsakloglou (2010). “The distributional impact of in-kind public benefits in European countries”. In: *Journal of Policy Analysis and Management* 29.2, pp. 243–266.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). “Distributional National Accounts: Methods and Estimates for the United States”. In: *The Quarterly Journal of Economics* 133.2, pp. 553–609.

Piketty, Thomas, Li Yang, and Gabriel Zucman (2019). “Capital Accumulation, Private Property, and Rising Inequality in China, 1978-2015”. In: *American Economic Review* 109.7, pp. 2469–2496.

Pinkovskiy, Maxim and Xavier Sala-i-Martin (2016). “Lights, Camera … Income! Illuminating the National Accounts-Household Surveys Debate”. In: *The Quarterly Journal of Economics* 131.2, pp. 579–631.

Sala-i-Martin, Xavier (2006). “The World Distribution of Income: Falling Poverty and... Convergence, Period”. In: *The Quarterly Journal of Economics* 121.2, pp. 351–397.

Soares, Sergei (2019). “The Market Value of Public Education - A Comparison of Three Valuation Methods”. CEQ Working Paper 71.

Statistics Canada (2019). “Distributions of Household Economic Accounts, estimates of asset, liability and net worth distributions, 2010 to 2018, technical methodology and quality report”. In.

Stiglitz, Joseph, Amartya Sen, and Jean-Paul Fitoussi (2009). *Report by the Commission on the Measurement of Economic Performance and Social Progress*. Tech. rep. INSEE.

van der Weide, Roy, Christoph Lakner, Daniel Gerszon Mahler, Ambar Narayan, and Rakesh Gupta (2024). “Intergenerational Mobility around the World: A new database”. In: *Journal of Development Economics* 166.1, pp. 103–167.

Verbist, Gerlinde, Michael Förster, and Maria Vaalavuo (2012). “The Impact of Publicly Provided Services on the Distribution of Resources: Review of New Results and Methods”. OECD Working Paper.

Wagstaff, Adam, Marcel Bilger, Leander R. Buisman, and Caryn Bredenkamp (2014). “Who Benefits from Government Health Spending and Why?” World Bank Policy Research Working Paper 7044.
World Bank (2013). *Measuring the Real Size of the World Economy: The Framework, Methodology, and Results of the International Comparison Program – ICP*. The World Bank.
– (2018). *Poverty and Shared Prosperity 2018: Piecing Together the Poverty Puzzle*. Tech. rep.

Figure 1 – Public Services and Global Poverty Reduction:
Real Average Income of the World's Poorest 20%, 1980-2022

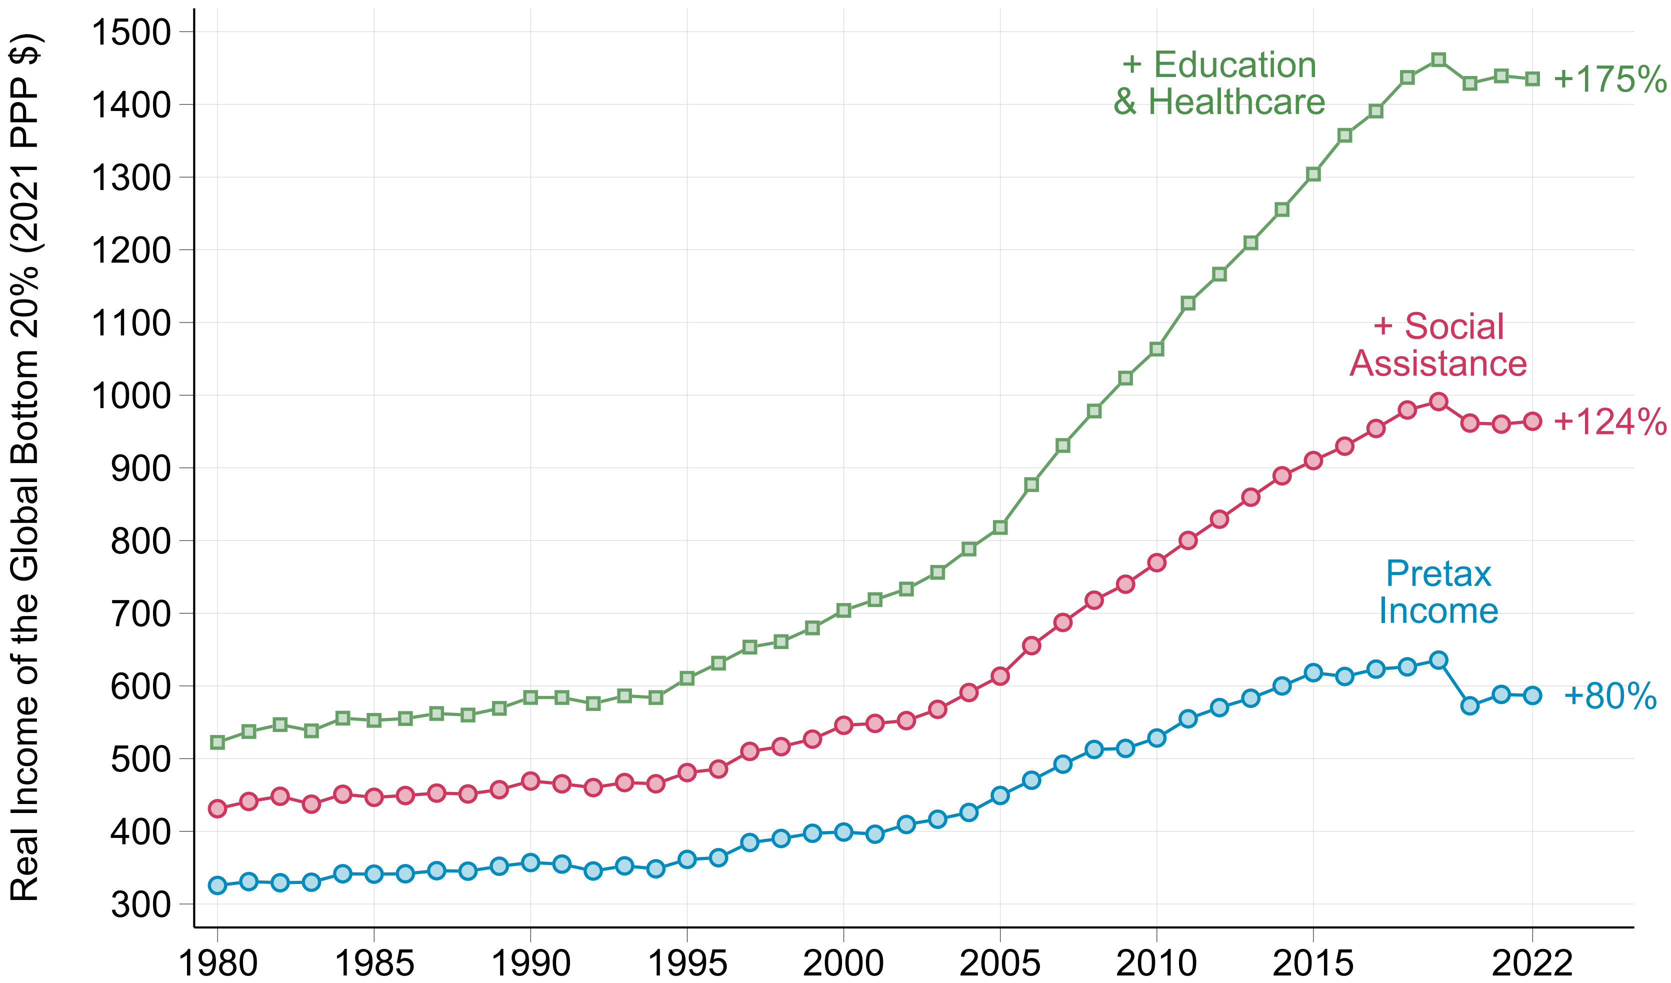

Notes. The figure plots the evolution of the world's poorest 20% real average income from 1980 to 2022, before and after adding social assistance transfers and the consumption of public education and healthcare to individual incomes. Pretax income grew by 80%, pretax income plus social assistance grew by 125%, and pretax income plus all government transfers grew by 175%. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare.

Figure 2 – Public Services and Poverty Measurement

(a) Public and Private Services are Substitutes: The Case of Healthcare

(b) Public Services Matter for Non-Monetary Dimensions of Quality of Life

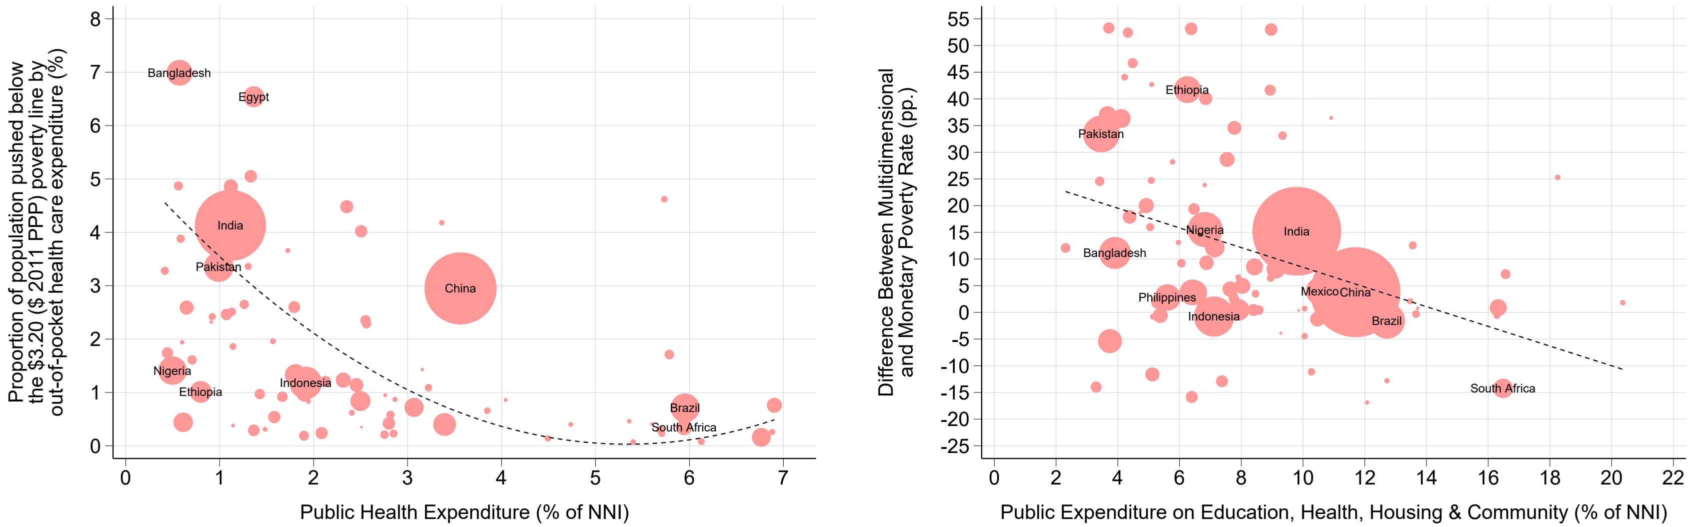

Notes. Panel A: Author's computations combining national budget data (public health expenditure) and World Bank estimates (healthcare-driven poverty). The figure plots the relationship across countries between public health spending, expressed as a share of national income, and healthcare-driven poverty, measured as the share of the population falling into poverty due to out-of-pocket health expenditure. In countries spending more on public healthcare, fewer households fall into poverty due to own spending on healthcare. Panel B: Author's computations combining national budget data (public expenditure), World Bank estimates (monetary poverty rate), and Oxford Poverty and Human Development Initiative estimates (multidimensional poverty rate). The figure plots the relationship across countries between public expenditure on education, health, housing, and community services, and the gap between monetary and multidimensional poverty measures. Monetary poverty: share of population whose consumption falls below $2.15 per day (2017 PPP USD). Multidimensional poverty: index combining deprivation in health, education, and living standards (see Alkire, Kanagaratnam, and Suppa, 2021). In countries with greater spending on basic public services, fewer households fall into multidimensional poverty relative to those falling in monetary poverty.

Figure 3 – Validation of Methodology: Comparison With High-Quality Studies.
Transfers Received by the Bottom 50%, This Paper versus Detailed Studies

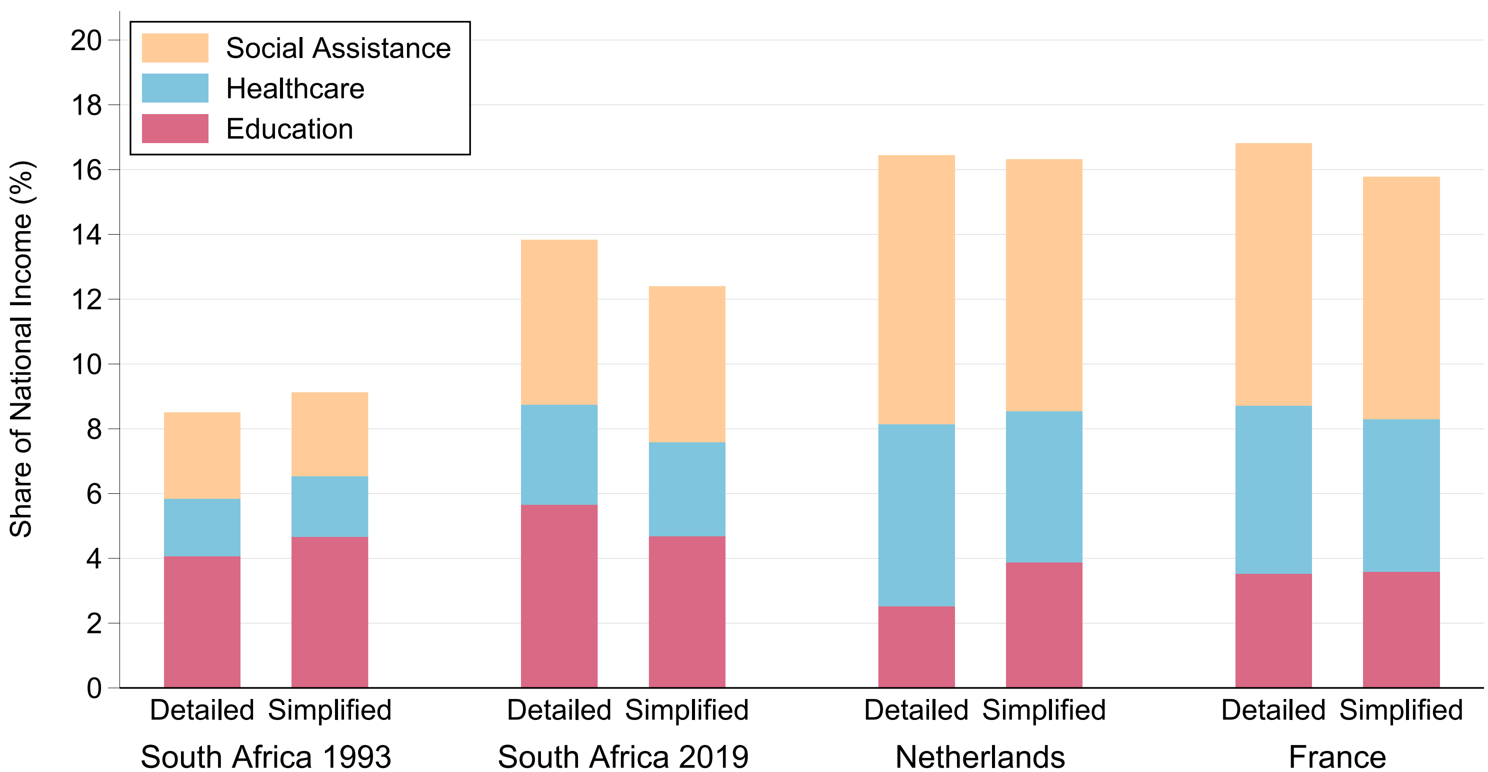

Notes. The figure plots the level and composition of government transfers received by the bottom 50% in France, the Netherlands, and South Africa, expressed as a percentage of national income, comparing simplified series (this paper) to detailed series constructed in Germain et al. (2021), Bruil et al. (2022), and Gethin (2024), respectively.

Figure 4 – Validation of Methodology: Comparison with CEQ and Over Time

(a) Comparison With CEQ: Education

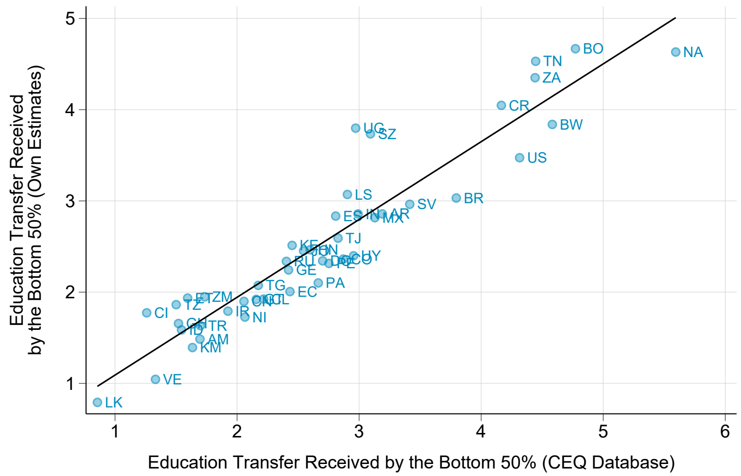

(b) Comparison With CEQ: Healthcare

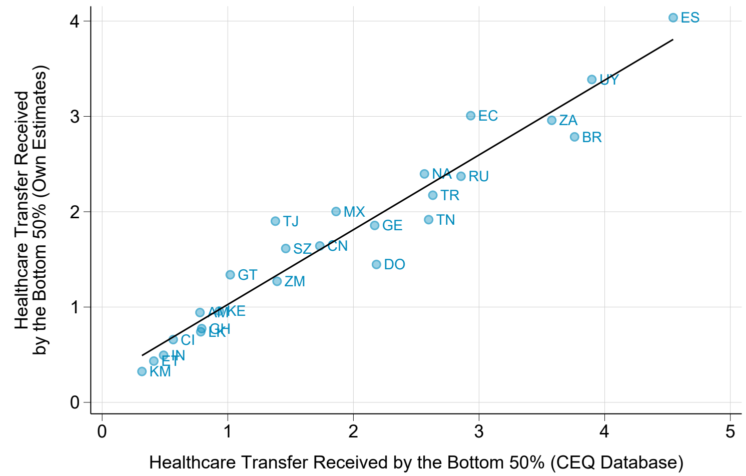

(c) Trends in Progressivity: Education

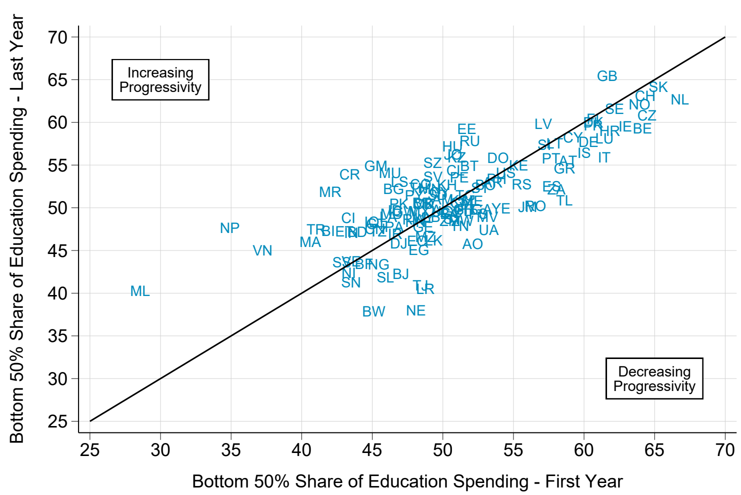

(d) Trends in Progressivity: Healthcare

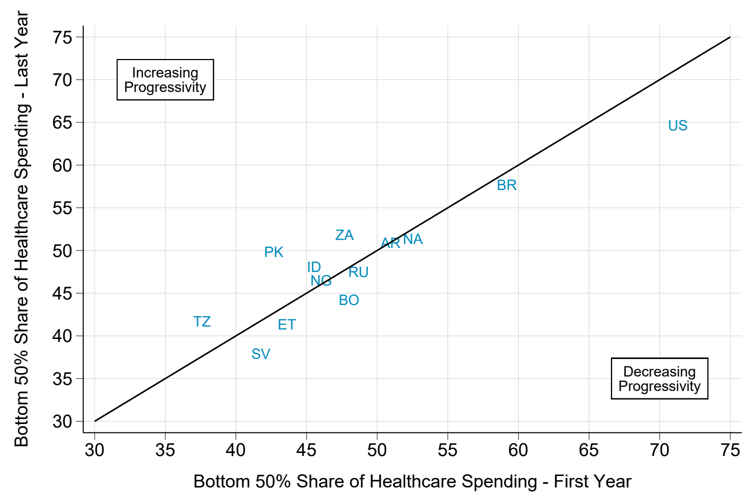

Notes. Panels (a) and (b) compare estimates of public education and health transfers received by the bottom 50% developed in this paper to high-quality estimates available in the Commitment to Equity (CEQ) Database. Panels (c) and (d) compare the progressivity of public education and health transfers in the first and last years available in each country.

Figure 5 – Public Spending on Education, Healthcare, and
Social Assistance by Country Income Group, 1980-2022

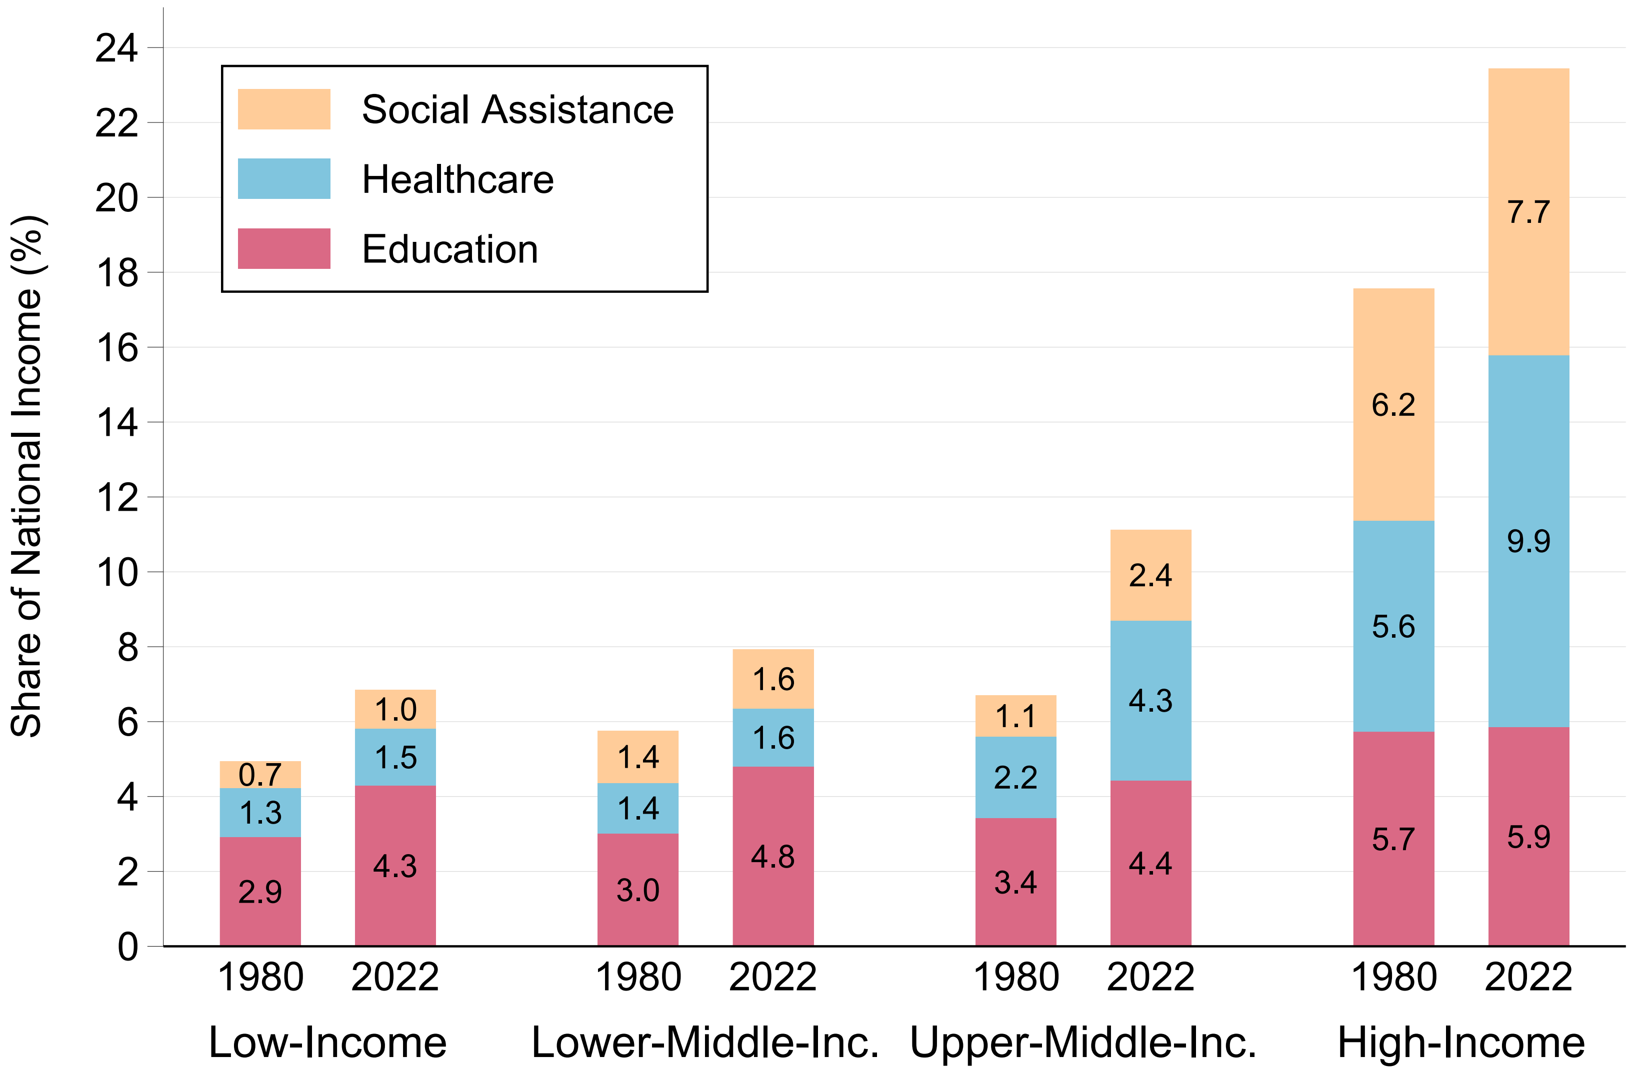

Notes. The figure plots the average share of national income spent on social assistance, public education, and healthcare by country income group. Public spending on these three categories of government transfers increased as a share of national income in all groups. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare. Population-weighted averages across all countries in each group.

Figure 6 – Share of Pretax Income and Government Transfers
Received by the Bottom 50% by Country Income Group

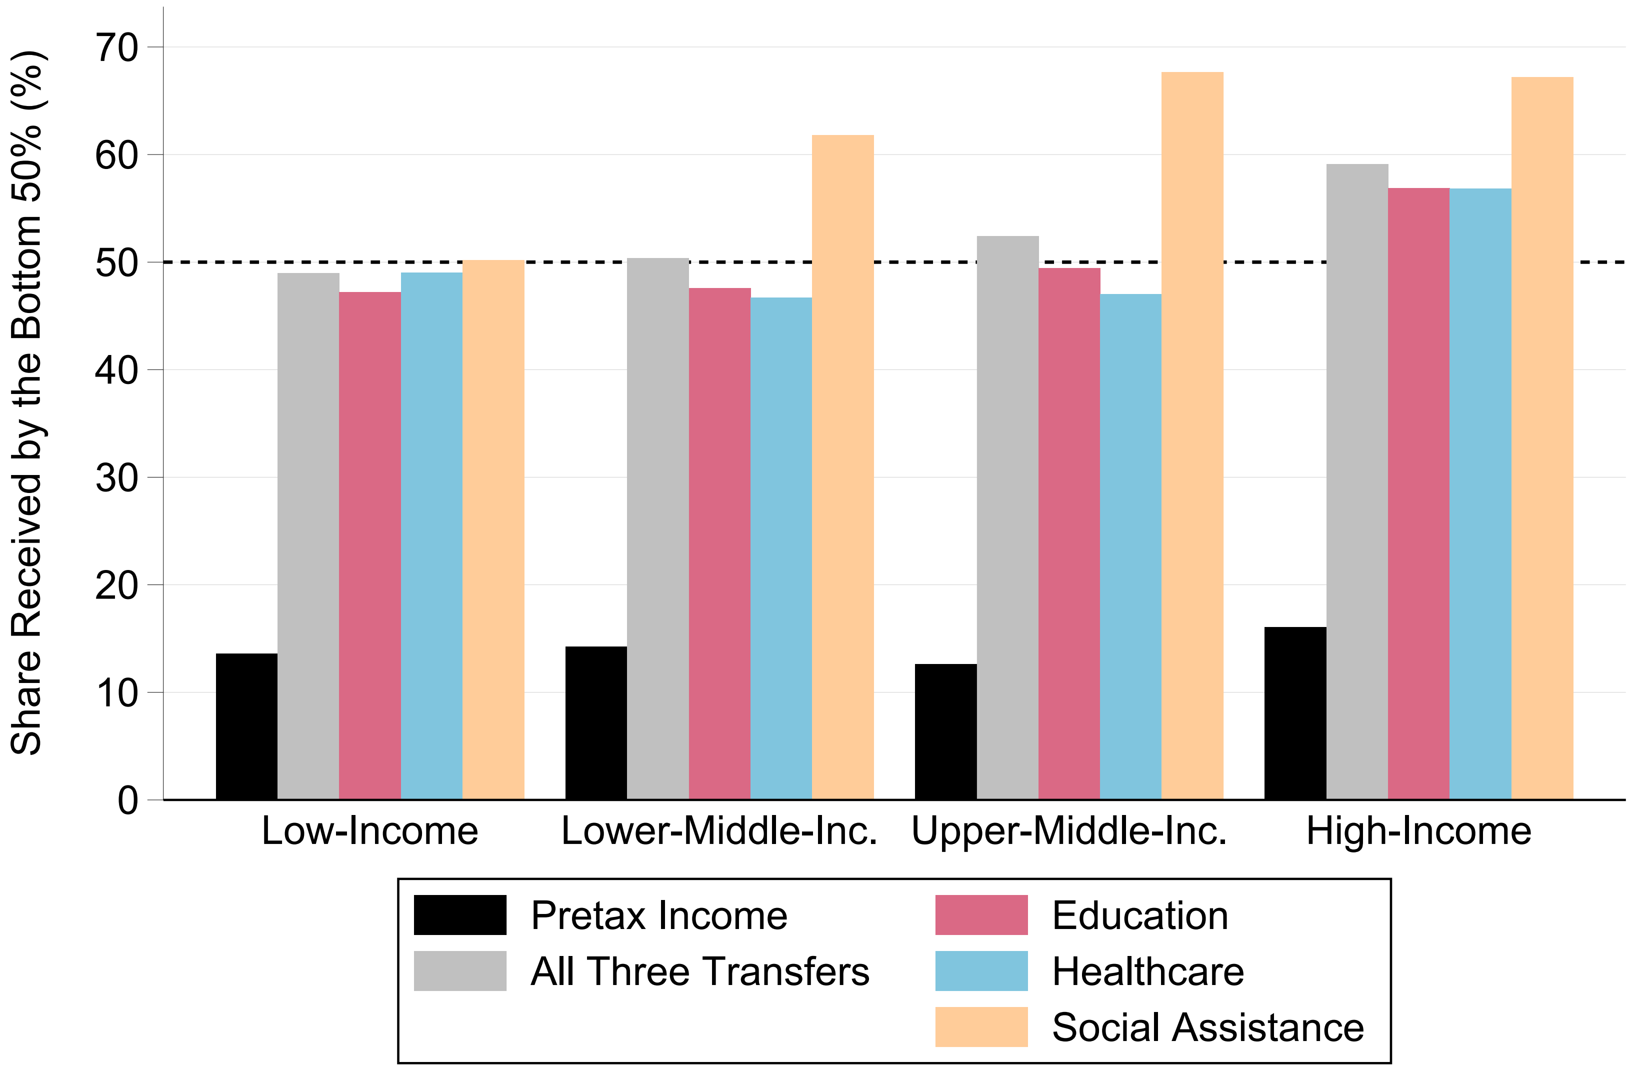

Notes. The figure plots the average share of pretax income, social assistance, public education, and healthcare received by the bottom 50% by country income group. All government transfers reduce inequality: they are less concentrated than pretax income. Transfers are more progressive in high-income countries than in low-income countries: the poorest 50% receive a greater proportion of each type of transfer. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare. All three transfers: the sum of social assistance, public education, and healthcare. Population-weighted averages across all countries in each group.

Figure 7 – Government Redistribution in International Perspective

(a) Social Assistance, Education, and Healthcare Transfers
Received by the Bottom 50% (% of National Income)

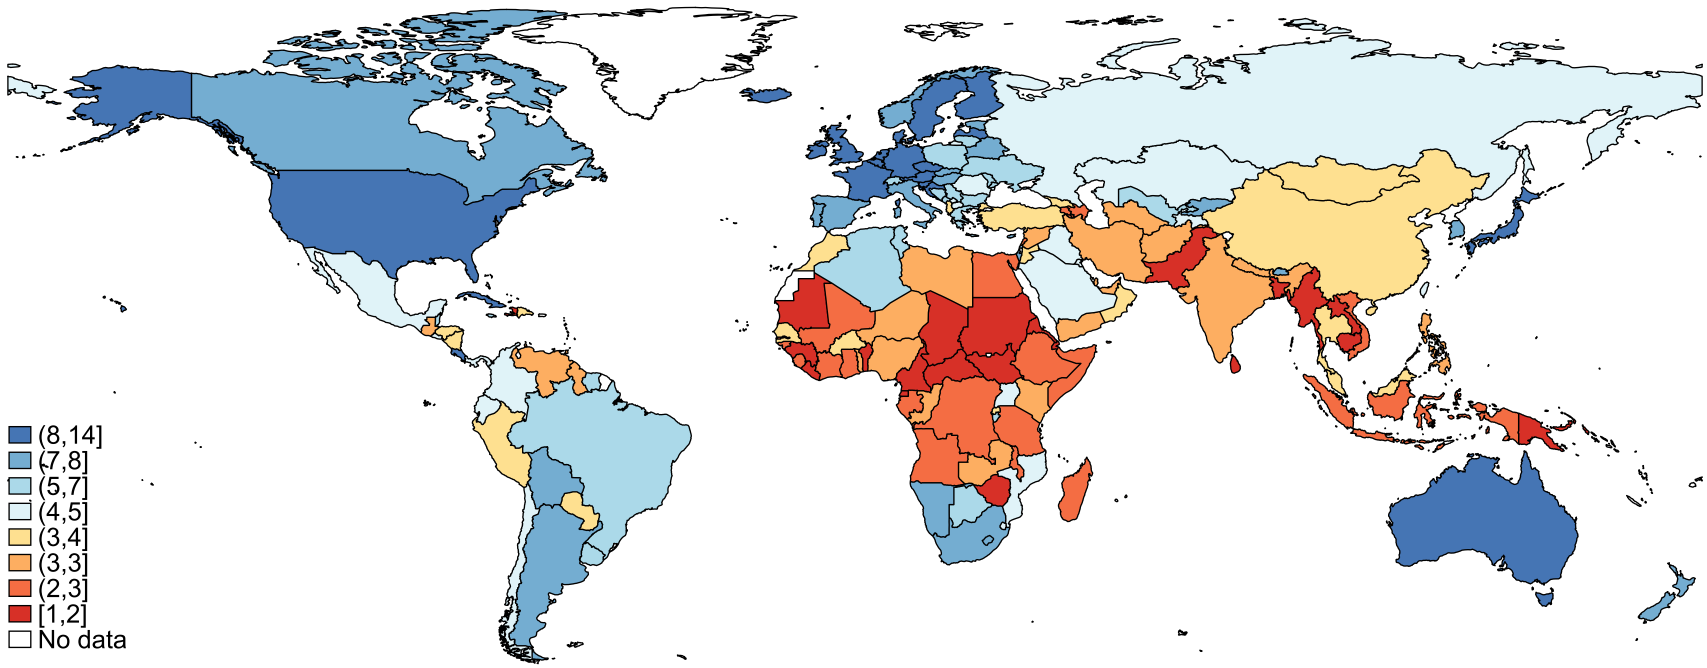

Notes. The figure maps total government transfers (social assistance, education, and healthcare) received by the bottom 50% in each country in 2022, expressed as a share of national income. Government redistribution is substantially larger in Western Europe and Northern America than in South Asia and Sub-Saharan Africa. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare.

Figure 7 – Government Redistribution in International Perspective

(b) Government Transfers Received by the Bottom 50% in Selected Countries

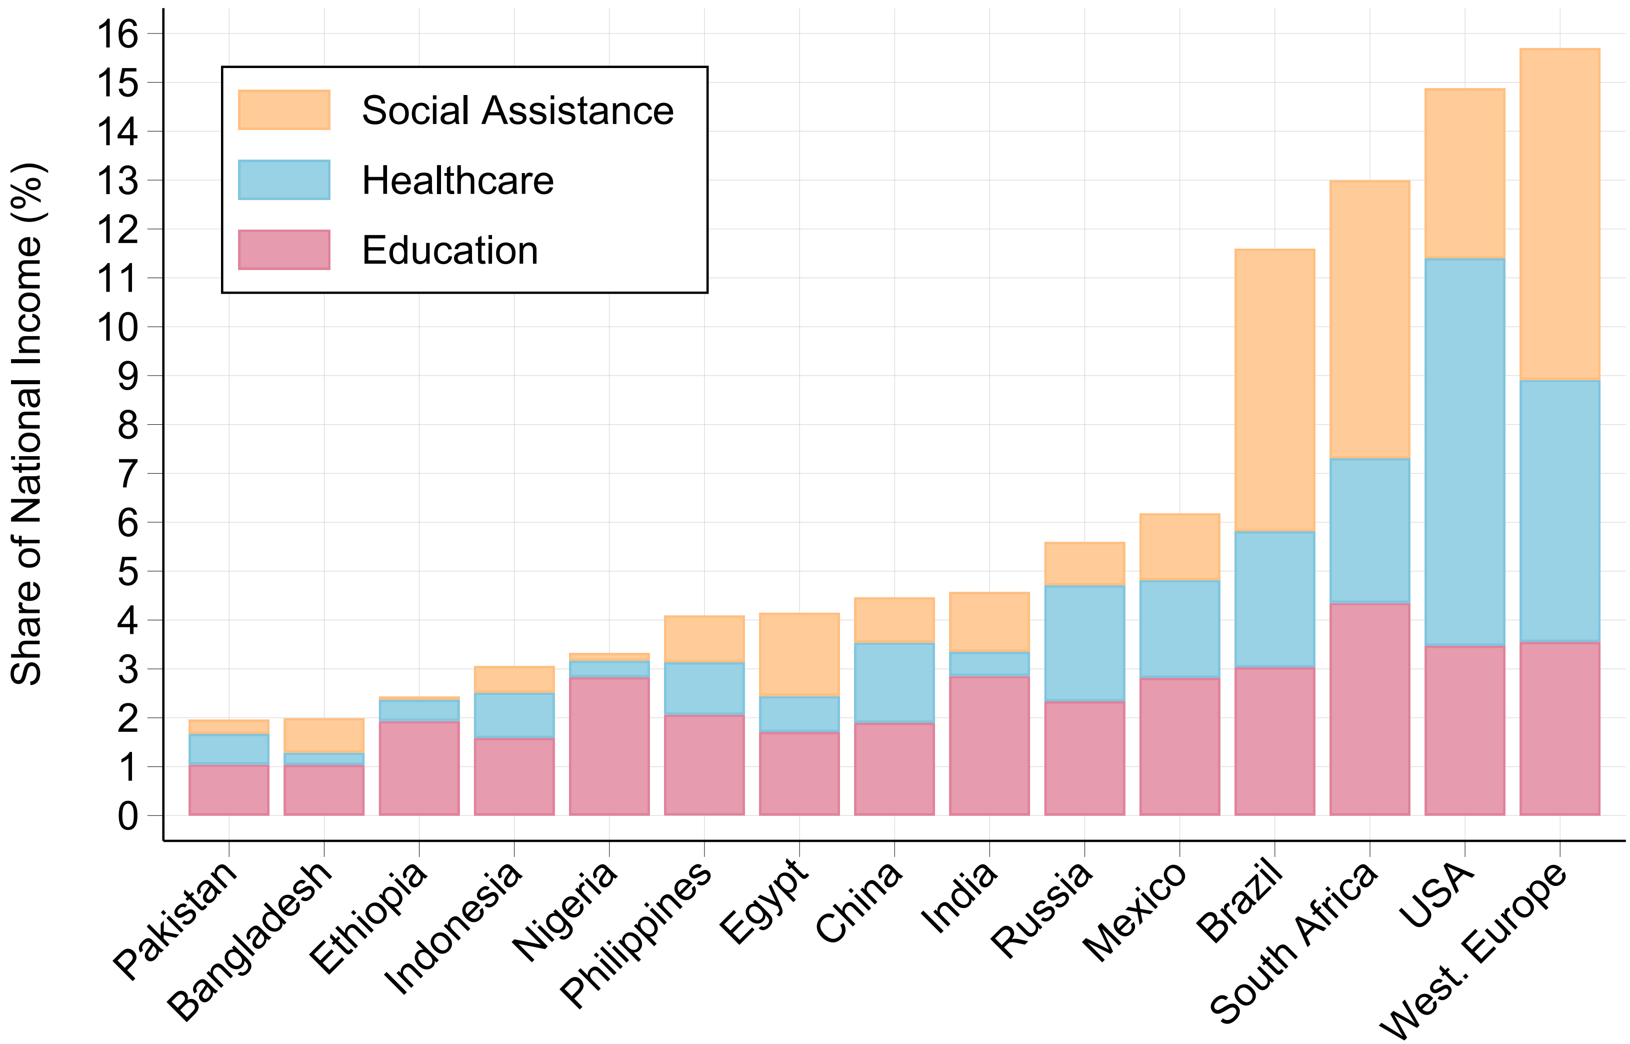

Notes. The figure plots the level and composition of government transfers received by the bottom 50% in selected countries or world regions in 2022, expressed as a share of national income. The poorest 50% receive almost 16% of national income in the form of social assistance, education, and healthcare in Western Europe, compared to only 2% in Pakistan. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare.

Figure 8 – Rising Government Redistribution Around the World:
Government Transfers Received by the Bottom 50%, 1980 versus 2022 (% of National Income)

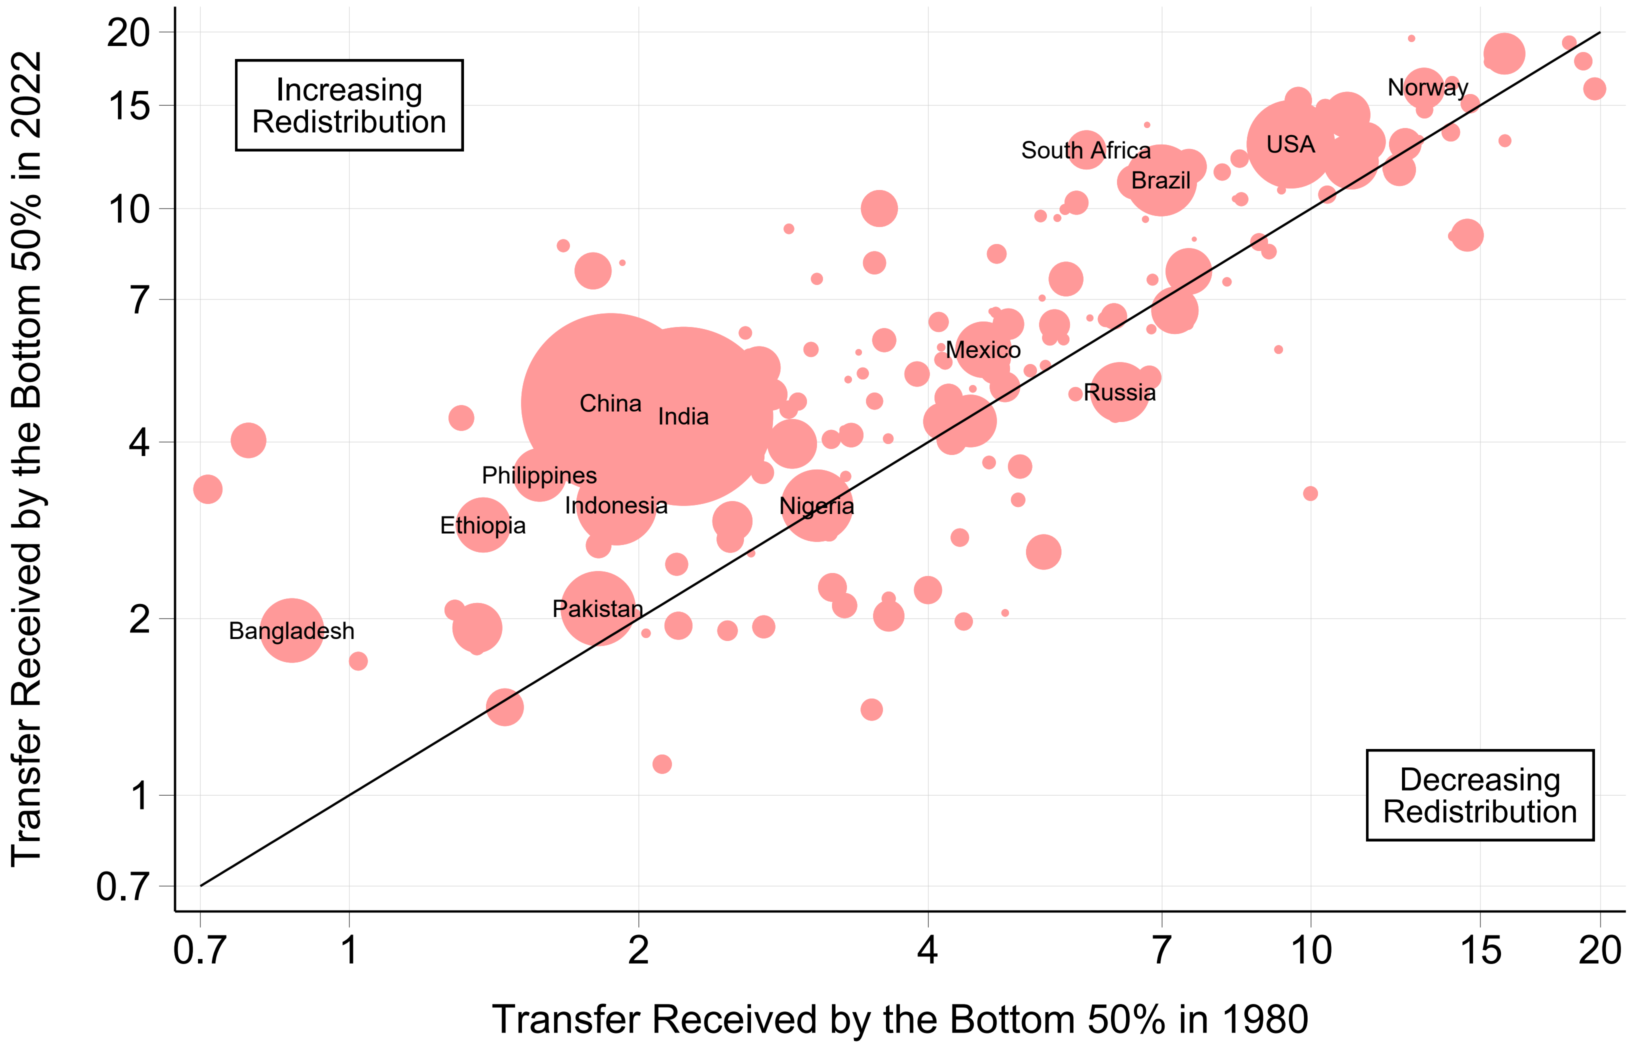

Notes. The figure compares government transfers received by the bottom 50% in each country in 1980 (x-axis) and 2022 (y-axis), expressed as a share of national income. Countries above the 45-degree line saw redistribution increase, while those below the 45-degree line saw redistribution decrease during this period. Government transfers include social assistance, public education, and healthcare. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare.

Figure 9 – Government Transfers Received by the World's Poorest 20%, 1980-2022

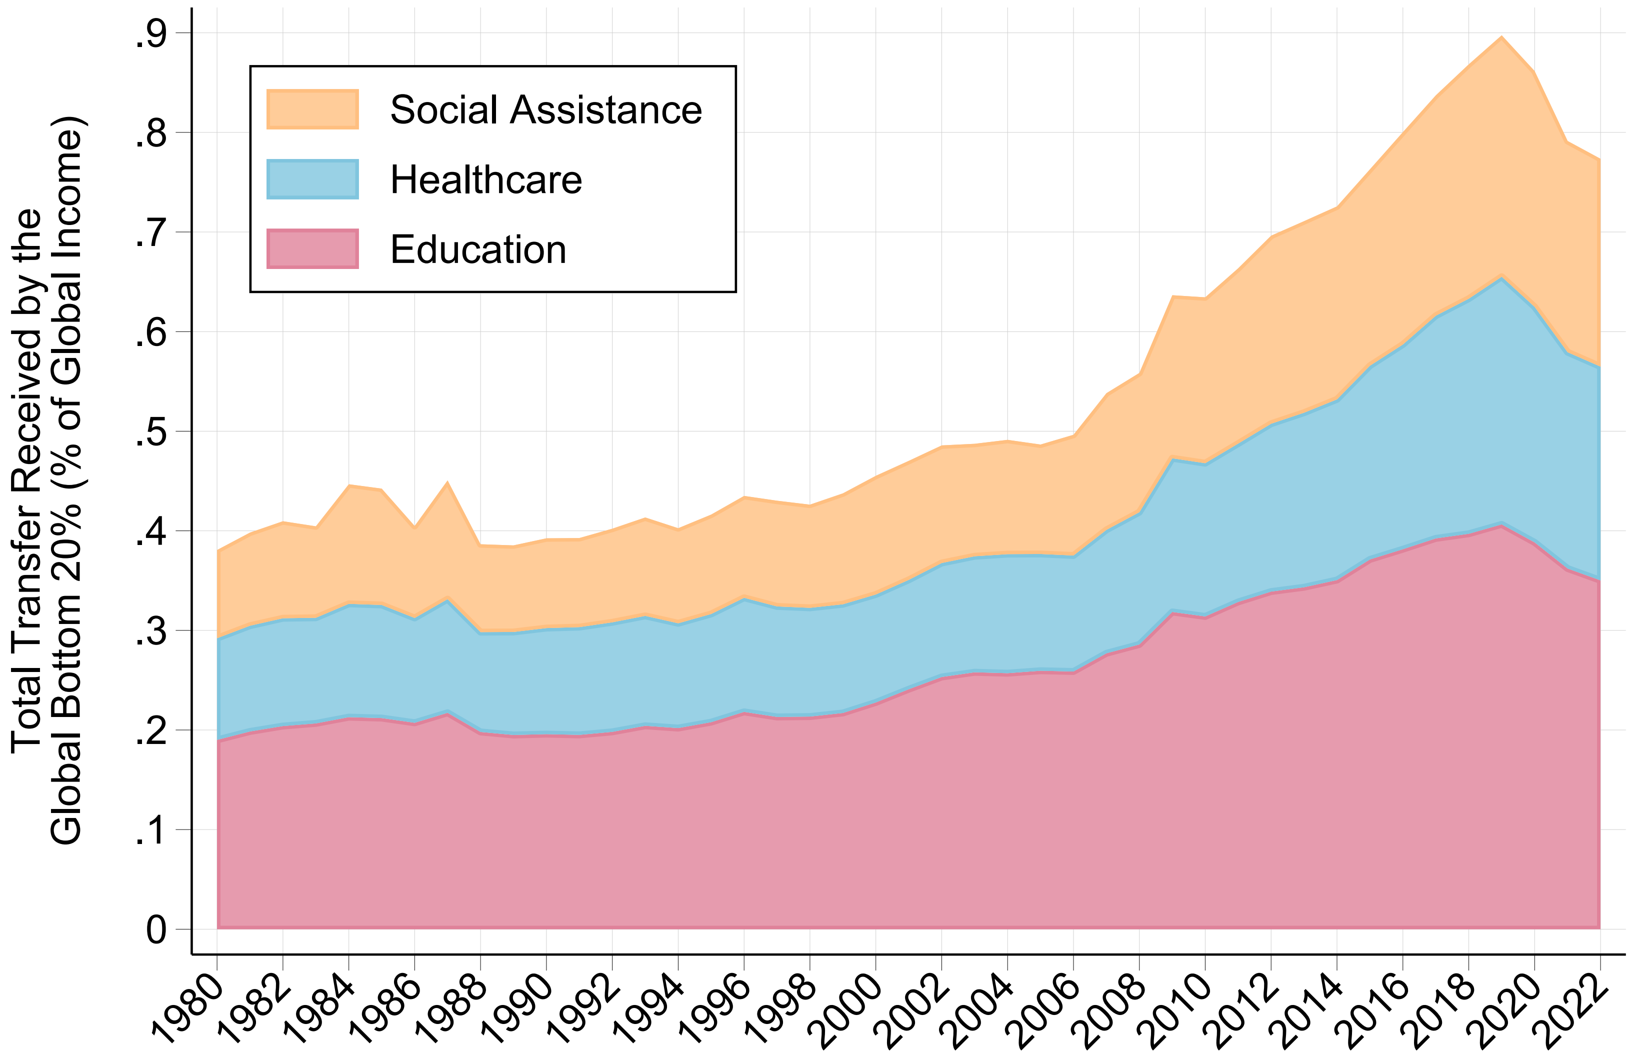

Notes. The figure plots the level and composition of government transfers received by the world's poorest 20% from 1980 to 2022, expressed as a share of global income (the sum of net national incomes across all countries in the world). Total transfers received by the global poor in the form of social assistance, education, and healthcare expanded considerably during this period, from about 0.4% to 0.8% of global income. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare.

Figure 10 – Government Redistribution and the World Distribution of Income:
Real Income Growth Rate by Global Income Percentile, 1980-2022

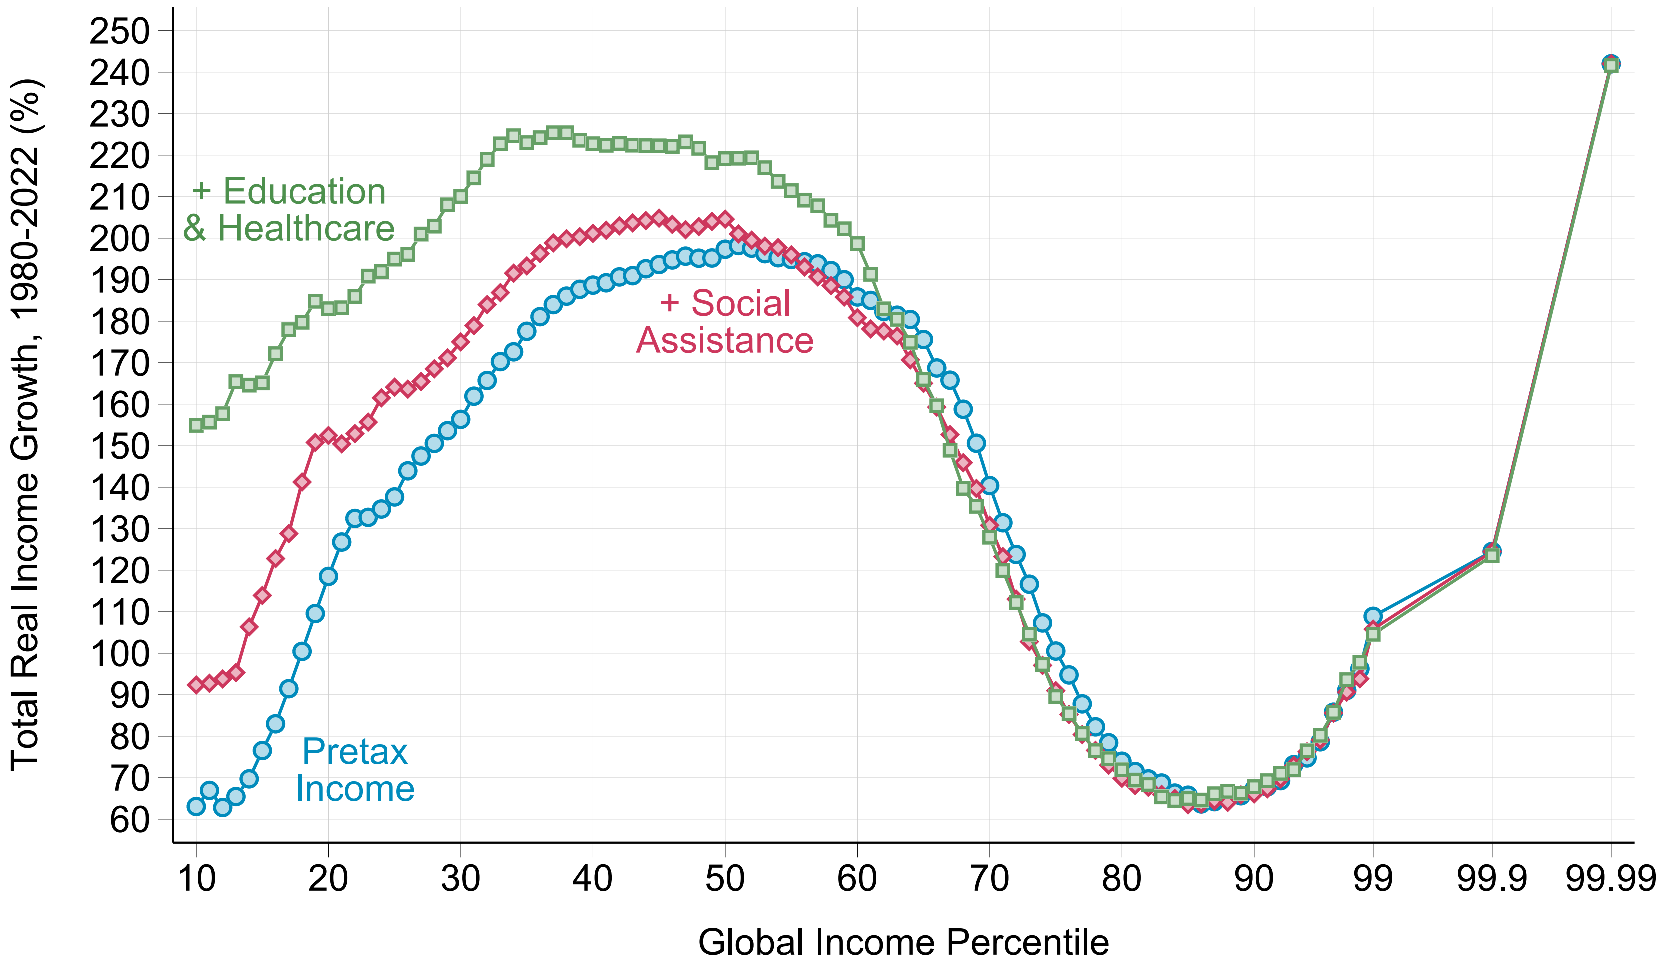

Notes. The figure plots total real income growth by global income percentile from 1980 to 2022, before and after adding social assistance transfers and the consumption of public services to individual incomes. Government transfers have played a key role in enhancing real income growth at the bottom of the world distribution of income. The average income of the 30th percentile grew by 160% in terms of pretax income, compared to 210% after adding social assistance, education, and healthcare transfers to individual incomes. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare.

Figure 11 – Accounting for Government Productivity

(a) Education Spending and Expected Human Capital

(b) Health Spending and Quality of Healthcare

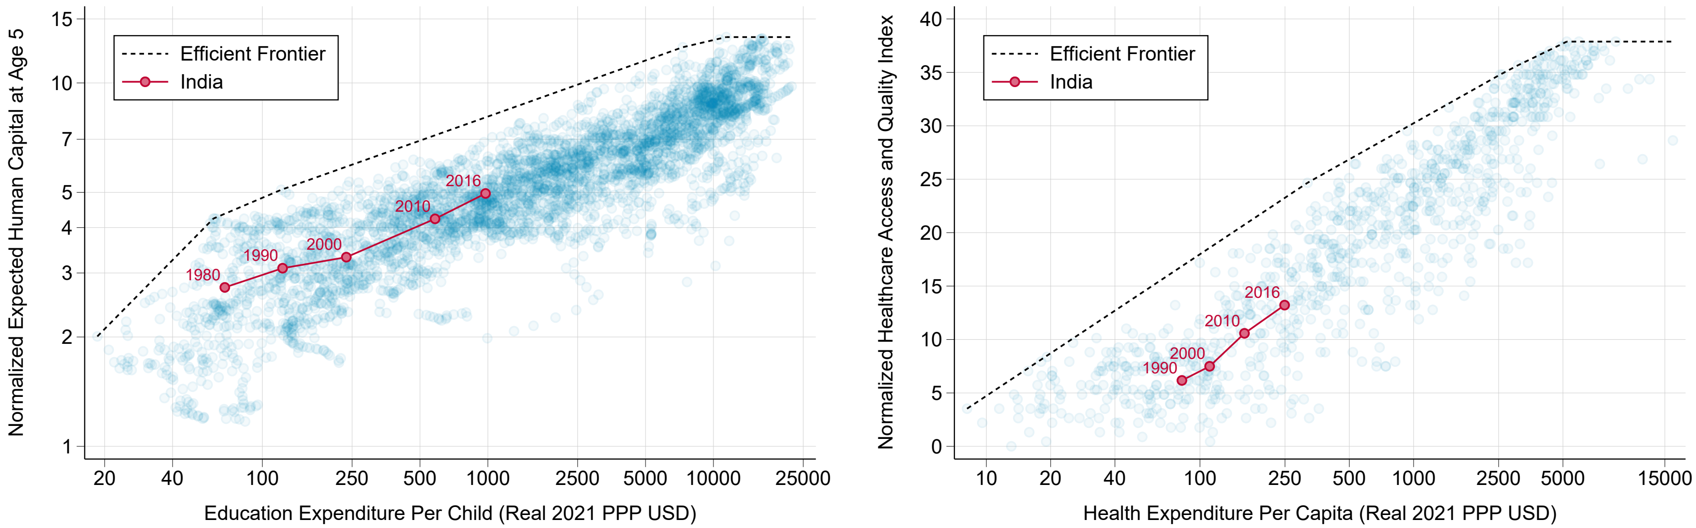

Notes. Panel A: the figure plots the relationship between education expenditure per child and expected human capital from age 5 onward. Panel B: the figure plots the relationship between healthcare expenditure per capita and the quality of healthcare, measured as the Healthcare Access and Quality index reexpressed in units of life expectancy (GBD, 2022). Both panels: each data point is a country-year. The upper dashed line represents the efficient frontier, defined as a piecewise linear estimate of the maximum achievable output at each level of expenditure. The trajectory of India is highlighted, revealing a significant increase in both public spending and the quality of education and healthcare received in this country in the past decades. India remains far from the frontier in both cases, however, revealing low levels of productivity in this country (that is, low education and healthcare outcomes given their cost of provision).

Figure 12 – Cost of Provision vs. Welfare Value of Public Services

(a) Education

(b) Healthcare

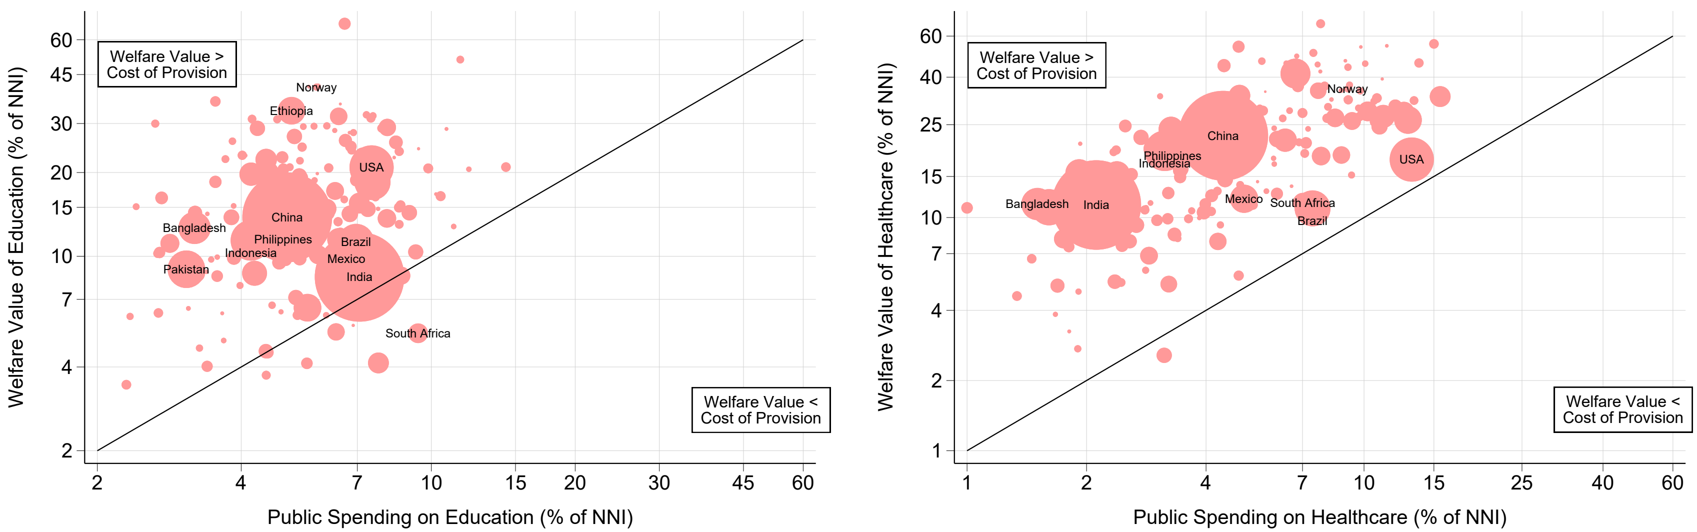

Notes. The figure compares total public education (panel a) and public healthcare (panel b) transfers received in each country, expressed as a share of national income, depending on whether they are valued at cost of provision (x-axis) or at their net present value (y-axis). The welfare value of public education and healthcare is higher than its cost of provision for countries above the 45-degree line; it is below cost of provision for countries below the 45-degree line. In most countries, the welfare value of education and healthcare significantly exceeds their cost of provision.

Table 1 – A New Database on Government Redistribution

	Number of Countries Covered		Avg. Share of NNI (%)		Bottom 50% Share (%)		
	Aggregates G	Distribution γ	1980	2022	Min	Mean	Max
Education	173	145	3.7%	4.8%	38%	49%	65%
Health	173	109	2.6%	3.7%	29%	49%	66%
Social Assistance	167	130	2.2%	2.7%	16%	64%	92%
All Three Transfers			8.5%	11.3%	35%	52%	66%
Pretax Income	173	173	100%	100%	5%	14%	28%

Notes. This table reports information on the geographical coverage of the database, together with selected descriptive statistics. The second column reports the number of countries covered in terms of aggregate public spending by function and total pretax income (G). The third column reports the number of countries covered in terms of pretax income inequality and the distributional incidence of transfers (γ). Columns 4 and 5 report the share of national income that each component represented in the average country in 1980 and 2022 (population-weighted average in each year). Columns, 6, 7, and 8 provide information on the minimum, maximum, and average share of each component that is received by the bottom 50%, calculated across all country-years with available data. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision in columns 4 and 5, that is, as total general government expenditure on education and healthcare. All three transfers: the sum of social assistance, public education, and healthcare.

Table 2 – Government Redistribution and the Geography of Global Inequality

	Pretax Income	+ Social Assistance	+ Education & Healthcare
Theil Decomposition			
Theil Index	1.18	1.12	1.02
Between-Country Component	30%	33%	38%
Within-Country Component	70%	67%	62%
Share in Global Bottom 20%			
India	17%	19%	21%
China	11%	11%	8%
Pakistan	17%	21%	27%
Bangladesh	15%	17%	19%
Ethiopia	40%	53%	63%
Nigeria	21%	30%	33%
Other Asia / Oceania	16%	16%	16%
Other Sub-Saharan Africa	60%	66%	69%
Middle East and Northern Africa	16%	16%	15%
Latin America	24%	17%	13%
US / Canada / Western Europe	8%	2%	0%

Notes. The table reports a Theil decomposition of global inequality into a between-country and a within-country component, as well as the geographical composition of the global bottom 20% in 2022, for different income concepts. Column 3 adds social assistance to pretax incomes, while column 4 adds both social assistance and the consumption of public services to pretax incomes. In 2022, the Theil index of global inequality was 1.18 in terms of pretax income, compared to 1.02 after adding all government transfers. 17% of the Indian population belonged to the world's poorest 20% in terms of pretax income, compared to 21% after transfers. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare.

Table 3 – Public Services, Quality of Life, and the GDP-Survey Gap

	Youth Literacy		Secondary School Enrollment Rate		Infant Mortality		Life Expectancy	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: No FE								
GDP-Survey Gap	0.07***	-0.01	0.28***	0.09***	0.42***	-0.17***	0.06***	-0.00
	(0.02)	(0.02)	(0.03)	(0.02)	(0.06)	(0.03)	(0.01)	(0.00)
Educ./Health Spending		0.11***		0.30***		0.59***		0.06***
		(0.01)		(0.01)		(0.01)		(0.00)
Panel B: Country FE								
GDP-Survey Gap	0.06**	0.04*	0.12***	-0.09***	0.36***	-0.01	0.04***	0.00
	(0.02)	(0.03)	(0.03)	(0.03)	(0.05)	(0.03)	(0.01)	(0.01)
Educ./Health Spending		0.02***		0.31***		0.67***		0.06***
		(0.01)		(0.01)		(0.01)		(0.00)
N	351	351	1412	1412	1762	1762	1762	1762
Adj. R-squared	0.83	0.84	0.81	0.88	0.89	0.95	0.88	0.92

Notes. Each column presents coefficients of a regression of a selected dependent variable on the gap between GDP and survey means (as in Pinkovskiy and Sala-i-Martin, 2016), before and after controlling for public spending on education or healthcare. GDP-Survey Gap: percentage difference between GDP per capita and survey mean income. Educ./Health Spending: log of public education spending (youth literacy, secondary school enrollment rate) or log of public healthcare spending (infant mortality, life expectancy). Panel A runs simple OLS regressions. Panel B includes country fixed effects. Countries with growing gaps between GDP and survey aggregates saw greater improvements in all four indicators of quality of life. However, this effect disappears in most specifications when controlling for public spending on education and healthcare. In other words, public services can rationalize the totality of the growing inability of surveys at tracking the evolution of living standards.

Revisiting Global Poverty Reduction: Public Services and the World Distribution of Income, 1980-2022

Amory Gethin

August 2024

SUPPLEMENTARY ONLINE APPENDIX

This appendix supplements my article "Revisiting Global Poverty Reduction: Public Services and the World Distribution of Income, 1980-2022." It contains additional methodological details, as well as supplementary figures and tables.

Table of Contents

A. Additional Methodological Details

This section presents the methodology used to estimate a new database on the level and composition of general government expenditure from 1980 to 2022. Section A.1.1 presents data sources. Section A.1.2 presents the methodology used to harmonize total general government expenditure series. Sections A.1.3, A.1.4 A.1.5, A.1.6, and A.1.7 turn to the harmonization of education, health, social protection, other components of government expenditure, and debt service cost, respectively. Section A.1.8 covers general government revenue.

A.1. General Government Expenditure Data

A.1.1. Data Sources

To construct comparable estimates of general government revenue and expenditure, I collect data from a number of available sources.

Data on general government expenditure and its composition by function of government come from eight sources:

1. IMF historical government revenue and expenditure series, based on initial data compilation by Mauro et al. (2015) and updates.[1] The data cover total government revenue, total government expenditure, interest paid on public debt, and government primary balance for over 150 countries from 1800 to 2022 (unbalanced panel).

2. IMF series of government expenditure by function.[2]. The data cover total expenditure and its breakdown by function (COFOG) for different levels of government, including general government, central government, state government, and local government series.

[1] https://www.imf.org/external/datamapper/datasets/FPP

[2] https://data.imf.org/?sk=388dfa60-1d26-4ade-b505-a05a558d9a42&sid=1479329334655

3. IMF series of government expenditure by function from historical government finance statistics. This data source is comparable to the previous one but covers earlier years.[3]

4. Eurostat series of government expenditure by function, covering the European Union.[4]

5. OECD series of government expenditure by function, covering OECD countries.[5]

6. CEPAL series of government expenditure by function, covering Latin American countries.[6]

7. The SPEED database (Statistics on public expenditures for economic development) maintained by the International Food Policy Research Institute (IFPRI), compiling series of government expenditure by function from various sources.[7]

8. United Nations series of government final consumption expenditure by function.[8]

Additional data on education and health expenditure come from:

1. The UNESCO UIS statistics, covering series of public education expenditure by level (primary, secondary, tertiary).[9]

2. The UNESCO's World Education Report 1991, which provides additional data points for 1980 and 1988 in a number of countries.

3. Bharti and Yang (2023), covering China and India since 1980.

4. The World Health Organization, covering health expenditure in almost all countries in the world since 2000 (available from the World Bank's World Development Indicators).

Additional data on social protection expenditure come from:

[3] https://www.icpsr.umich.edu/web/ICPSR/studies/8624/publications
[4] https://ec.europa.eu/eurostat/data/database
[5] https://stats.oecd.org/Index.aspx?DataSetCode=SNA_TABLE11
[6] https://statistics.cepal.org
[7] https://www.ifpri.org/publication/statistics-public-expenditures-economic-development-speed
[8] http://data.un.org/Data.aspx
[9] http://data.uis.unesco.org/

1. The Asian Development Bank, covering total social protection expenditure in a number of Asian countries.[10]

2. The CEPAL SOCX database, covering social protection expenditure and its decomposition by type of program in Latin American countries.[11]

3. The OECD SOCX database, covering social protection expenditure and its decomposition by type of program in OECD countries.[12]

4. The World Bank SPEED database, covering social protection expenditure and its decomposition by type of program in selected Eastern European and Central Asian countries.[13].

5. The World Bank ASPIRE database, covering social protection expenditure and its decomposition by type of program in a number of developing countries.[14].

A.1.2. General Government Expenditure

I start by constructing harmonized series of total general government expenditure for all countries. I do so in two steps.

First, I attribute a "benchmark series" to each country, corresponding to the primary data source used to observed the level of total expenditure. I assign one benchmark source to each country, with the following order of priority: OECD, CEPAL, IMF historical series, other IMF series (general government when available, or the sum of central and local governments, or central government), SPEED, and the World Bank. For countries not covered by any of these sources, I assume that expenditure equals total general government tax revenue, as reported in Bachas et al. (2022). For countries still not covered, I use regional averages of expenditure as a share of GDP.

[10]https://spi.adb.org/spidmz/
[11]https://statistics.cepal.org
[12]https://stats.oecd.org/Index.aspx?datasetcode=SOCX_AGG
[13]https://www.worldbank.org/en/programs/speed
[14]https://www.worldbank.org/en/data/datatopics/aspire/indicator-glance

Second, I use other sources to extend benchmark series backwards and forwards. For instance, if the OECD covers the 2000-2022 period in a given country while the IMF covers 1980-2022, I extrapolate OECD series backwards using absolute changes in expenditure as a percentage of GDP available from the IMF over the 1980-2000 period. This ensures that harmonized series do not have sudden ups and downs due to inconsistencies across sources, while taking the best of available information on annual changes in expenditure from each source. I follow the same order of priority as above. As a last resort, to go back to 1980, I assume that total expenditure has remained constant as a share of GDP when no other information is available.

A.1.3. Education

Given that education spending series have greater coverage than series providing the full breakdown of expenditure by function of government, I choose to separately estimate education spending rather than directly estimating the full decomposition of government expenditure. The same principle holds for healthcare and social protection. I thus construct series of education, health, and social protection spending separately, and then combine them with other information on the composition of expenditure to estimate other functions of government as a residual.

To estimate harmonized public education expenditure series, I follow the same principles as those used for total expenditure.

First, I attribute a benchmark series by order of priority: general government education expenditure from Eurostat, the OECD, CEPAL, the IMF, and the UNESCO, followed by central government expenditure from the same sources if no general government expenditure series is available. Data for China and India are taken from Bharti and Yang (2023).

Second, I carry these benchmark series backward and forward using other available sources, with the same order or priority. To cover the full 1980-2022 period, I assume that education expenditure has remained constant as a share of total general government primary expenditure when no other information is available.

Third, I estimate the breakdown of education expenditure by level (primary, secondary, tertiary). I combine all available data points from UNESCO, CEPAL, and the IMF, by order of priority. Data for China and India are taken from Bharti and Yang (2023). To cover the full 1980-2022 period, I assume that the composition has remained constant when no other information is available.

A.1.4. Healthcare

I follow similar steps to estimate harmonized series of public health expenditure. The benchmark series are general government health expenditure from Eurostat, the OECD, CEPAL, the IMF, the WHO, and SPEED, by order of priority. In addition to WHO series available online, which cover the 2000-2022 period, I add historical series from Murray, Govindaraj, and Musgrove (1994), which report estimates of public health expenditure in 153 countries in 1990. I carry these series backward and forward using sources other than the benchmark, as for education. To cover the full 1980-2022 period, I assume that health expenditure has remained constant as a share of total general government primary expenditure when no other information is available.

A.1.5. Social Protection

I follow similar steps to estimate harmonized series of social protection expenditure.

First, I construct series of total social protection expenditure, including social insurance and social assistance programs. The benchmark series are general government social protection expenditure from Eurostat, the OECD, CEPAL, and the IMF, followed by central government expenditure from the same sources if no general government expenditure series is available. For countries with no data from any of these sources, I use estimates from the World Bank's ASPIRE database, the United Nations' final consumption expenditure series, the SPEED database, or the

Asian Development Bank, depending on which of these sources is available and seems most consistent. I carry these series backward and forward using sources other than the benchmark, as for education and health. To cover the full 1980-2022 period, I assume that social protection expenditure has remained constant as a share of total general government primary expenditure when no other information is available.

Second, I split social protection expenditure into social insurance and social assistance. For consistency, I only use one source by country: the OECD SOCX database, the CEPAL SOCX database, and the World Bank ASPIRE database, by order of priority. For the 41 countries with data from none of these sources, I use regional averages across countries with available data. To cover the full 1980-2022 period, I assume that this share has remained constant when no data is available.

Third, I split social assistance expenditure into cash and in-kind transfers. This breakdown is available from the OECD SOCX database and the World Bank ASPIRE database, covering together 127 countries. In the remaining countries, I use regional averages across countries with available data. To cover the full 1980-2022 period, I assume that this share has remained constant when no data is available.

A.1.6. Other Components of Government Primary Expenditure

The other components of government primary expenditure are expenditure on general public services (excluding debt service cost), defense, public order and safety, economic affairs, environment protection, housing and community amenities, and recreation, culture and religion. Data on these components of expenditure are much scarcer and much more difficult to link across sources. I estimate them in two steps.

First, I construct series on the composition of expenditure covering all functions of government. Given difficulties in linking series, I assign only one source by country: general government expenditure series from Eurostat, the OECD, CEPAL, and the IMF, by order of priority, followed by central government expenditure series (or the sum of local and central governments) from the same sources when general government series are not available. For countries with no data from any of these sources, I use regional averages of the composition of spending by COFOG.

Second, I ensure consistency between these series and the total expenditure, health expenditure, education expenditure, and social protection expenditure series estimated separately above. I assume that these four series are correct and proportionally adjust the relative importance of other functions of government. This amounts to estimating the size of these other functions of government as a residual (total expenditure minus education, health, and social protection), while using the above sources to measure their importance relative to one another. In the rare cases where this leads to negative value (that is, the sum of education, health, and social protection exceeds total expenditure), I assume that the education, health, and social protection series are correct, that other functions of government represent 1% of total expenditure each, and I adjust total expenditure series upwards accordingly.

A.1.7. Debt Service Cost

I estimate debt service cost series separately, using the same methodology as for education, health, and social protection expenditure. The benchmark series are historical IMF series, general government expenditure IMF series, and central government expenditure IMF series, by order of priority. To cover the full 1980-2022 period, I assume that debt service cost has remained constant as a share of total general government expenditure when no other information is available. For missing countries, I use regional averages of debt service cost as a share of GDP.

Having estimated debt service cost, I calculate general government primary expenditure, equal to total general government expenditure minus debt service cost.

A.1.8. General Government Revenue

Finally, I estimate the level and composition of general government revenue using the same methodology. The benchmark series are historical IMF series, which I carry backward and forward using total tax revenue estimates from Bachas et al. (2022). In the few cases where total tax revenue from Bachas et al. (2022) exceeds IMF series, I assume that total government revenue equals tax revenue. I then calculate non-tax revenue as the difference between general government revenue and tax revenue. The composition of tax revenue by type of tax directly comes from Bachas et al. (2022), including personal income taxes, corporate income taxes, property and wealth taxes, indirect taxes, other taxes, and social contributions. To cover the full 1980-2022 period, I assume that government revenue as a share of GDP and its composition have remained constant when no data is available. In missing countries, I use regional averages of government revenue as a share of GDP.

A.2. Distributional Incidence Data

A.2.1. Education

To estimate the distributional incidence of public education, I construct a new survey microdatabase on inequality in access to schooling worldwide. The database covers nationally representative information on the following variables: household and individual identifiers, age, gender, school attendance, and household consumption or disposable income. The surveys are assembled from three sets of sources.

The first data source is the World Bank's I2D2 database. Since the 1990s, the World Bank has made considerable effort to assemble and harmonize nearly all living standards surveys ever fielded in the world (such as the CPS in the U.S., PNAD in Brazil, or SUSENAS in Indonesia). The resulting database, which I was able to access at the World Bank, contains approximately a thousand surveys fielded in 156 surveys over the 1980-2012 period.

The second data source is the World Bank's Global Monitoring Database. It is the direct successor of I2D2, which was discontinued at the end of the 2010s. It compiles over 4,200 microfiles, bringing together surveys fielded in 170 countries over the 1980-2021 period. There is some overlap between the I2D2 and GMD, although imperfect (some historical surveys are in I2D2 but not in GMD). When both sources cover a given country-year, I give priority to GMD.

The third data source consists in country-specific surveys, which I use to complement the I2D2 and GMD (and in some cases to improve the coding of variables in surveys already covered by these two sources). I collect the raw survey data files from international or country-specific data portals, and harmonize them to cover the variables outlined above. I was able to do so for a total of 51 countries covered over the 1981-2019 period. These include Brazil, China, the Democratic Republic of the Congo, Egypt, Ethiopia, Ghana, Indonesia, India, Iraq, Jordan, Mexico, Nepal, Nigeria, Pakistan, Palestine, Russia, South Africa, Tanzania, Thailand, the United States, and the 31 countries covered by the European Union Statistics on Income and Living Conditions survey (EU-SILC). Table B8 provides information on the resulting country and time coverage of the microdatabase.

Drawing on this database, I calculate the transfer received by per capita disposable income or consumption group p in country c at time t as:

$$g_{pct}^{\text{educ}} = n_{pct}^{pri} g_{pct}^{pri} + n_{pct}^{sec} g_{pct}^{sec} + n_{pct}^{ter} g_{pct}^{ter} \tag{8}$$

Where n_{pct}^{k} denotes the average number of children in school at level k, g_{pct}^{k} denotes average spending per child on function k, and $k \in \{pri, sec, ter\}$ refers to primary education, secondary education, and tertiary education, respectively.

School attendance by level and income decile is directly estimated from the microdatabase. In the absence of data on the exact grade followed by each child, I allocate primary education spending to children in school who are aged 0 to 12, secondary education to those aged 12 to 18, and tertiary education to those older than 18. I then extrapolate decile-level cells of school attendance to the entire 1980-2022 period by assuming that relative intensity of use by decile and level has remained constant.

Data on public education expenditure by level come from the database compiled in this paper. I thus allocate spending by level proportionally to relative intensity of use of the education system by decile. This yields a total education transfer received by decile in each country-year.

A.2.2. Healthcare

To estimate the distributional incidence of healthcare, I rely on surveys reporting information on individuals' use of the healthcare system. I combine data from three sources.

The first source is the CEQ database. This database compiles data from country-specific studies following a comparable methodology to estimate the distributional incidence of taxes and transfers. Most of these studies allocate public healthcare expenditure by relying on surveys providing information on household members' frequency of use of the healthcare system (typically the number of visits to a doctor in the past month: see Lustig, 2018). This indicator is then aggregated by per-capita household disposable income decile. I use this database to cover 20 countries.[15].

The second data source consists in country-specific surveys covering similar information on actual use of the healthcare system. I collect and harmonize surveys from country-specific data portals for the Democratic Republic of the Congo, Nigeria, Pakistan, and Russia. These surveys ask household members whether they had a medical problem in the past month, and whether they saw a healthcare provider when they did. The question is slightly different for Russia, asking about the total number of visits to a doctor in the past year. I complement these surveys with the 2016 Life in Transition Survey (LITS), covering 34 Eastern European and Central Asian countries, which provides information on whether respondents received medical care in the

[15] Although this database covers 45 countries, I prefer to use the other two data sources described thereafter to cover the remaining 25. Figure 4b compares estimated measures of the distributional incidence of healthcare in the CEQ to those estimated using the other two sources.

past 12 months.

The third data source is the World Health Survey, which was fielded by the WHO around 2003. Similarly to other surveys, the WHS asks whether the household member needed healthcare in the past month, and whether they did receive healthcare when they did. I use this survey to cover 53 additional low- and middle-income countries.

Combining these sources, I calculate the total number of visits made to a healthcare provider by per-capita household income or consumption decile. I then distribute public healthcare expenditure proportionally to this measure of intensity of use of the healthcare system.

A.2.3. Social Assistance

To distribute social assistance expenditure, I combine data from four sources.

Data on the United States come from Piketty, Saez, and Zucman (2018) and updates. It covers the entire 1980-2022 period.

Data on 30 European countries come from Blanchet, Chancel, and Gethin (2022). Estimates are based on household surveys—the Luxembourg Income Study or EU-SILC depending on the country. They cover the year 2019.

Data from the World Bank's ASPIRE database allows me to cover 96 additional countries. The data covers benefits incidence curves for total social assistance transfers received by income or consumption quintile. Most data points come from surveys fielded in the 2010s.

Data from the CEQ database allows me to cover four remaining countries. The database reports information on the share of social assistance transfers received by decile.

These four data sources allow me to cover the distributional incidence of social assistance for a total of 130 countries. I thus distribute total social assistance expenditure proportionally to these incidence profiles in each country-year.

A.2.4. Taxes

The distributional incidence of taxes come from a companion paper (Fisher-Post and Gethin, 2023). In this paper, we combine aggregate data on the composition of tax revenue by type of tax with assumptions on the distributional incidence of each type of tax in 151 countries since 1980. The methodology follows the Distributional National Acccounts principles (Blanchet et al., 2021). In broad strokes, personal income taxes are distributed using a specific tax simulator, accounting for the differential taxation of labor and capital income. Corporate income taxes are distributed proportionally to corporate equity. Property and wealth taxes are distributed proportionally to housing wealth (residential property taxes) and capital income (business property taxes and inheritance and wealth taxes). Indirect taxes are distributed proportionally to consumption. See Fisher-Post and Gethin (2023) for more details.

A.2.5. World Bank Income Distribution Data

As a robustness check, I investigate the sensitivity of my results to using World Bank income distribution data instead of the World Inequality Database. The data is available from the World Bank's Poverty and Inequality Platform (PIP), and presents itself in the form of distributions available for selected countries and years. The income concept is either consumption or posttax disposable income per capita, depending on the country. All values are reported in 2017 PPP USD per day.

Unfortunately, although the World Bank regularly reports indicators of global poverty, it does not publish underlying estimates of the world distribution of income. I thus attempt to reconstruct measures of pretax and posttax income myself. Starting from available data, I first extrapolate the average income of each country-percentile to missing years using real GDP per capita growth rates. For the 17 countries entirely missing, I use estimates from the World Inequality Database. The resulting database yields trends in global poverty almost identical to those officially reported by the World Bank. I then reconstruct measures of pretax income and posttax income. In the absence of any information on savings, I define pretax income in each country as consumption or disposable income, minus social assistance transfers, plus direct taxes.

Both the levels and trends in global poverty in the WID data differ from those of the World Bank for at least four main reasons. First, World Bank estimates mostly focus on consumption (posttax disposable income minus net household saving), while my focus here is on income. The difference between consumption and income can be large, with major implications on levels and trends in inequality in some cases (see for instance Chancel et al., 2023). Second, the estimates presented here are consistent with national income growth rates, while World Bank estimates are based on surveys and do not attempt to bridge gaps between survey and national accounts aggregates. Third, some of the estimates used in this paper are based on detailed country-specific studies that rely on data sources that may differ from those of the World Bank in a number of countries, including China (Piketty, Yang, and Zucman, 2019), India (Chancel and Piketty, 2019), and Latin America (De Rosa, Flores, and Morgan, 2022). See Chancel and Piketty (2021). Fourth, I use GDP purchasing power parity conversion factors, while the World Bank only corrects for price differences in household final consumption expenditure. For all these reasons, I view the World Inequality Database as a more adequate source for conducting the particular analysis developed in this paper. A more systematic comparison of these two datasets is left to future research.

A.3. Public Sector Productivity

A.3.1. Methodology

As explained in section 5, I estimate public sector productivity by benchmarking the productivity of governments around the world to one another. If a government produces more output than any other for a given cost, then its efficiency is set to 1, and the productivity of other governments with comparable costs is estimated based on the outputs they deliver.

I estimate public sector productivity by combining data on education and healthcare expenditure (inputs) and outcomes (outputs). I first choose a given input (e.g., education expenditure) and output (e.g., human capital). I then use data envelopment analysis to non-parametrically estimate the technical frontier, defined as the maximum output ever achieved in any country-year for a given level of expenditure and other inputs (e.g., Herrera and Ouedraogo, 2018). Finally, I use the estimated frontier to estimate Θ^j, based on the extent to which output could be improved without changing costs in a given country-year.[16] This yields measures of technical efficiency ranging from 0 to 1 for each country-year covered by the data.

I apply this methodology to separately estimate the productivity of public education and public healthcare. When data is not available for the whole 1980-2022 period, I extrapolate backwards and forwards by assuming that productivity has remained constant. For countries with no data at all, I take the global average observed in each year.

A.3.2. Education

Input The first element required to estimate public education productivity is a measure of cost of provision. I take public education expenditure per child, expressed in 2021 PPP US dollars, estimated from the public spending database compiled in this paper.

Output The second element needed is a measure of government performance. Following the large literature in macroeconomics investigating the role of education in explaining differences in economic development (e.g., Hanushek, Ruhose, and Woessman, 2017), I propose to measure the output of the education system as the expected human capital that a child can hope to obtain at age 5:

$$Y^{\text{education}} = \exp(r_S S + r_Q Q) \tag{9}$$

[16] I use the `teradial` command in Stata.

With S expected years of schooling at age 5, r_S the return to a year of schooling, Q a measure of education quality, and r_Q the return to education quality. Data on expected years of schooling come from the UNESCO and covers 202 countries over the 1970-2020 period. Education quality is taken from Altinok, Angrist, and Patrinos (2018), who compile data from various international test scores to construct a new database of education quality in 134 countries. The return to schooling is set to 10% per year and the return to quality to 15% per standard deviation, as is standard in the existing literature.

Results Figure 11a plots the resulting relationship between performance and cost of provision for all country-years. There is a very strong correlation between the two variables: countries spending more on education display education systems of substantially better quality. Yet, there is also significant dispersion in the expected human capital stock achieved for a given level of government expenditure. The upper dashed line represents the efficient frontier, estimated using data envelopment analysis with variable returns to scale. This corresponds to a piecewise linear estimate of the maximum achievable output by level of expenditure.[17]

A.3.3. Healthcare

Inputs As for public education, the first step is to collect data on cost of provision. Given the particular role that private healthcare can play in some countries, I focus on total healthcare expenditure per capita (private and public combined), available from the World Bank's World Development Indicators.

Output Finding an accurate measure of the quality of healthcare provision is more challenging than for education. Indeed, unlike the human capital stock, which has a clear cardinal

[17]Notice that as shown in figure 11a, I fit the efficiency frontier using the log of the human capital stock. To get correct efficiency measures, one then needs to convert the ratio of logs into the ratio of actual human capital stocks. More precisely, we have a measure of efficiency θ^{log} such that: $\theta^{log} = \frac{\log x}{\log \bar{f}(x)}$, with $\bar{x}$ the technical frontier evaluated at x. The objective is to convert θ^{log} into $\theta = \frac{x}{f(x)}$. Rearranging yields $f(x) = \exp(\frac{\log x}{\theta^{log}})$ and hence $\theta = \frac{x}{\exp(\frac{\log x}{\theta^{log}})}$.

(monetary) interpretation, there is no obvious measure of healthcare performance whose units are directly comparable to cost of provision. Quality-adjusted life expectancy is often taken as a measure of interest (e.g., Cutler et al., 2022), yet this indicator is, by itself, arguably a poor measure of the performance of the healthcare system. Given these limitations, I turn instead to the healthcare access and quality (HAQ) index estimated in the context of the global burden of disease study (GBD, 2022). This indicator ranks healthcare systems from 0 to 100, based on death rates from 32 causes of death that could be avoided by timely and effective medical care. The main advantage of the HAQ index is that it was specifically created by health experts to measure the ability of healthcare systems to cure preventable diseases: it is explicitly a measure of performance. It also has the advantage of covering nearly all countries in the world since 1990. The disadvantage is that it is normalized from 0 to 100, so it has no cardinal interpretation.

In the absence of better solution, I re-express the HAQ index in units of life expectancy by first regressing it on life expectancy at birth, and then normalizing it using the coefficient obtained. More specifically, I run a linear regression of life expectancy on the HAQ index, controlling for the log of GDP per capita, years of schooling of the working-age population, and country fixed effects. I then multiply the HAQ index by the coefficient obtained, so as to re-express it in "units of life expectancy." The resulting index ranges from less than 1 (Eritrea) to 34 (Finland), meaning that an average Finn would live 34 years less if there was no healthcare system at all.

Although existing estimates are scarce, it is useful to compare this indicator to findings from the existing literature. Cutler et al. (2022), in particular, use rich data on health outcomes in the United States to derive new measures of healthcare productivity growth over the 1999-2012 period. They find that medical treatment accounts for a total increase in quality-adjusted life expectancy of 1.7 years during this period. My own estimates using the transformed HAQ index indicate that life expectancy was higher by 36.6 years in the U.S. in 1999 thanks to the healthcare system, compared to 37.9 years in 2012. This represents 1.3 years increase in life expectancy thanks to medical treatment, slightly lower than the 1.7 years found by Cutler et al. (2022). This suggests that the transformed HAQ index provides a good approximation, and if anything might underestimate progress made in the quality of healthcare in the past decades.

Results Figure 11b plots the resulting relationship between healthcare performance and cost of provision for all country-years. As for education, there is a very strong correlation between the two variables: countries spending more on healthcare are much more able to limit deaths from curable diseases. The upper dashed line represents the efficient frontier.

A.3.4. Discussion: Estimates of Θ^j as Lower Bounds on Government Productivity

I view these estimates as providing a *lower bound* on government productivity, especially in poor countries, for three main reasons.

First, national income purchasing power parity conversion factors do already account for government productivity (World Bank, 2013). Indeed, public sector productivity is adjusted for all government services in the Asia-Pacific, Western Asia, and Africa regions, using a Cobb-Douglas function that assumes that government employees are less productive in poor countries because of a lower and less efficient stock of capital equipment (Heston, 2013). In OECD countries and the European Union, further adjustments are made for health and education, combining indicators on the quantity and quality of services provided (Blades, 2013). Hence, the correction made here to account for aggregate productivity implies adjusting transfers downwards twice, once when using PPP conversion factors to correct for price differences across countries, and once when multiplying transfers received by Θ^j.

Second, the frontier approach implies by construction that Θ^j cannot be greater than 1, given that the maximum input-output combination ever observed in any country-year is given a score of 1. As a result, governments are assumed to never be more productive than the private sector for any kind of service provided ($\Theta^j = 1$ corresponds to a government exactly as cost efficient as the private sector).

Third, omitted variable bias is likely to drive estimates of Θ^j in poor countries significantly *downwards*. Indeed, poor countries are likely to have worse outcomes for a given level of government expenditure not only because of inefficiencies, but also because of a number of other confounding factors. These include lower incomes, greater inequality, more extreme weather conditions, or lower human capital, which directly affect education and health outcomes independently from government investment. For all these reasons, overall government expenditure is likely to be more efficient in these countries than what the model suggests.

A.3.5. Validation: Correlates of Government Efficiency

Finally, a useful way of checking the reliability of my measures of government productivity is to compare them to existing indicators. Appendix Table A4 shows that education and healthcare productivity are positively correlated with a number of indicators of government efficiency available from international sources and the literature. This is especially true of healthcare productivity, which is positively associated with a composite index of government effectiveness ($\rho = 0.57$ for single-input estimates), lower corruption ($\rho = 0.43$), and more transparent policy-making ($\rho = 0.34$). I also find a positive correlation between my measures of healthcare efficiency and the index of public sector productivity of Chong et al. (2014) ($\rho = 0.29$), who mail letters to 159 countries and argue that the rate of return of these letters to their original sender provides a simple and transparent measure of government productivity.

A.4. Welfare Value of Public Services

A.4.1. Welfare Value of Education

Framework There are four levels of education s in the economy: no schooling (0), primary (1), secondary (2), and tertiary education (3). Individuals in country c can go to school for one year, delivering a return γ_{cs} for schooling level s when they start working.[18] They discount these future monetary benefits of education at a rate δ, and they receive them every year for 45 years of working life (say, from age 20 to 65).

The net present value of a year of primary education is:

$$NPV_{ctp1} = \sum_{i=t}^{t+45} \frac{(1+\gamma_{c1})y_{cip0} - y_{cip0}}{(1+\delta)^i} = \sum_{i=t}^{t+45} \frac{\gamma_{c1}y_{cip0}}{(1+\delta)^i}$$

In other words, an individual in primary school will get the return to a year of primary education (γ_{c1}), multiplied by the average income he would get if he had no schooling (y_{cip0}), summed up and discounted over the next 45 years ($i \in [t, t+45]$). I make the conservative assumption that being in school does not generate additional opportunities for reaching the next level (doing so would increase the net present value of education; see Soares, 2019).

Similarly, the net present values of secondary and tertiary education are:

$$NPV_{ctp2} = \sum_{i=t}^{t+45} \frac{\gamma_{c2}y_{cip1}}{(1+\delta)^i}$$

$$NPV_{ctp3} = \sum_{i=t}^{t+45} \frac{\gamma_{c3}y_{cip2}}{(1+\delta)^i}$$

An individual in secondary school benefits from the return to secondary education, multiplied by the average income of primary-educated workers. An individual in tertiary education benefits from the return to tertiary education, multiplied by the average income of secondary-educated workers.

Combining these formulas, the total transfer received by income group p is:

$$g_{ctp}^{\text{educ}} = n_{ctp1}NPV_{ctp1} + n_{ctp2}NPV_{ctp2} + n_{ctp3}NPV_{ctp3}$$

[18]For simplicity, the return to schooling is assumed to be constant over time and by income group. Empirical evidence suggests no systematic long-run trend in returns to schooling around the world: see Montenegro and Patrinos (2021). Evidence on heterogeneity in returns to schooling by income group is scarce, but suggests that if anything, returns are higher for low-income earners (e.g., Clay, Lingwall, and Stephens, 2021).

$$g^{\text{educ}}_{ctp} = \sum_{s=1}^{3} n_{ctps} \Bigg(\sum_{i=t}^{t+45} \frac{\gamma_{cs} y_{cip,s-1}}{(1+\delta)^i} \Bigg)$$

With n_{ctps} the number of children in country c at time t belonging to income group p enrolled in level s. Finally, let us assume that the economy grows at a constant rate g, so that $y_{cip,s-1} = (1+g)y_{c,i-1,p,s-1}$. The formula becomes:

$$g^{\text{educ}}_{ctp} = \sum_{s=1}^{3} n_{ctps} \Bigg(\sum_{i=t}^{t+45} \frac{(1+g)^i \gamma_{cs} y_{ctp,s-1}}{(1+\delta)^i} \Bigg)$$

Estimation I estimate this formula by combining data from several sources.

1. n_{ctps} is estimated as $n_{ctps} = n_{cts} \nu_{ctsp} \zeta^{\text{publ}}_{cts}$, with:
   (a) n_{cts} the number of pupils in country c at time t who are at level s. Data are available from the World Bank's World Development Indicators.
   (b) ν_{ctsp} the distribution of these pupils by income group. Data come from the microdatabase constructed in this paper.
   (c) $\zeta^{\text{publ}}_{cts}$ the fraction of pupils enrolled in public schools. Data are from the WDI.
2. γ_{cs} is taken from Gethin (forthcoming). In this paper, I construct a new microdatabase on education and individual incomes to produce comparable estimates of the returns to schooling by level in 150 countries.
3. y_{ctps} is the income that a given individual can expect to receive depending on his country c, the time period t, the income group of his household p, and schooling level s. I calculate it by combining data from the following sources:
   (a) Data on average pretax income by percentile for the 1980-2022 period come from the World Inequality Database.
   (b) Data on the expected income rank that an individual can hope to reach is estimated by using intergenerational mobility curves from van der Weide et al. (2024). For

instance, an average child in Brazil living in a family belonging to percentile 10 is expected to reach percentile 28 as an adult. This allows me to calculate, for each child, the income they can expect to receive as an adult depending on the income of their family of origin.

(c) Finally, I let this income vary by education level using estimates of relative incomes by education level from Gethin (forthcoming). For instance, an individual with no schooling receives 37% of the average income in Brazil. I thus assume that a child at percentile 10 can hope to receive 37% of the average income of percentile 28 as an adult. This gives me a complete mapping between the income group of a child attending school at level s and her expected income as an adult y_{ctps}.

4. δ is set at a standard value of 0.05.

Current Versus Future Income Specifications Ideally, one would like to have a complete mapping of both current and future incomes in order to estimate the net present value of education at a given point in time. For instance, a child ending school in 2022 is expected to receive a net benefit of $\gamma_{cs} y_{c,2023,p,s-1}$ in 2023, $\gamma_{cs} y_{c,2024,p,s-1}$ in 2024, etc. This highlights an important fact: the net present value of education is inherently dependent on future economic growth. Unfortunately, predictions over future growth are by nature uncertain, making it difficult to have an accurate measure of the true welfare value of schooling. I thus consider two alternative measures: one that relies on available predictions regarding the evolution of GDP in the coming decades, and one that assumes zero future economic growth ($g = 0$).

In the future income specification, which I take as the benchmark specification in the main body of the paper, I use actual predictions on the evolution of real GDP per capita in the coming decades. The OECD publishes GDP projections for 47 countries over the 2020-2060 period, covering about 60% of the world's population. For the remaining countries, I extrapolate GDP per capita to 2060 by assuming that real annual income growth will be similar to that observed over 2000-2022.[19] Notice that estimates of the welfare value of education in 1980 do not rely much on predictions, since the trajectory of incomes is already observed over 1980-2022. It is for the more recent years that estimates of real GDP per capita going up to 2050-2060 are needed. In all cases, I assume that inequality within countries will remain the same, meaning that the average income of each percentile grows at the same rate. With a complete dataset covering the distribution of income for each country-year over 1980-2060, I thus estimate the net present value of education for the full 1980-2022 period, accounting for future income growth.

In the current income specification, I make the very conservative assumption that the net present value of education is only dependent on current incomes. In other words, a child ending school in 2022 is expected to receive a net benefit of $\gamma_{cs} y_{c,2022,p,s-1}$ in 2023, $\gamma_{cs} y_{c,2022,p,s-1}$ in 2024, etc. The net present value of the education system becomes:

$$g_{ctp}^{\text{educ}} = \sum_{s=1}^{3} n_{ctps} \left(\sum_{i=t}^{t+45} \frac{\gamma_{cs} y_{ctp,s-1}}{(1+\delta)^i} \right)$$

A.4.2. Welfare Value of Healthcare

Framework In the same spirit, I estimate the welfare value of public healthcare based on the monetary value of additional years of life expectancy enabled by the healthcare system. Let V_{ctp} be the value of a year of life for percentile p in country c at time t. Individuals can expect to live T_{ct} years longer, L_{ct} years of which can be attributable to access to public healthcare. An average individual thus receives every year until their death an average annual life expectancy gain equal to $\frac{L_{ct}}{T_{ct}}$. The net present value of the public healthcare system is then:

$$g_{ctp}^{\text{heal}} = \zeta^{publ} \sum_{i=0}^{T_{ct}} \frac{\frac{V_{cip} L_{ct}}{T_{ct}}}{(1+\delta)^i} \tag{10}$$

[19] I bound annual growth rates between 1% and 5% to avoid extreme outliers.

ζ^{publ} is the fraction of life expectancy gains explained by public healthcare, proxied as the share of public spending in total healthcare expenditure. T_{ct} is calculated for simplicity as the difference between life expectancy and life expectancy divided by two (e.g., Hammitt and Robinson, 2011). The discount rate δ is set to 5%, as for education. For L_{ct}, the additional years of life that individuals benefit from thanks to the healthcare system, I rely again on the Healthcare Access and Quality index (HAQ) published by the Global Burden of Disease Study, which I reexpress in units of life expectancy at birth (see section A.3 for a longer discussion).

The remaining key parameter to estimate is V_{ctp}, the value of a statistical life year. Ideally, one would like to estimate V_{ctp} for every country and income group in the sample. Unfortunately, such detailed estimates are not available. I thus consider two alternative sets of specifications.

Discounted Income Specifications In my benchmark specification, I assume that the value of additional year of life only consists in the discounted future income flows that individuals can expect to receive. As discussed by Hammitt and Robinson (2011), this is a lower bound of the value of a statistical life year, since individuals arguably also value additional leisure from living longer. Additional income from living longer is only realized in $T_{ct} - L_{ct}$, when the individual starts benefiting from greater life expectancy, and is received until they die, in T_{ct}. The net present value of the healthcare system becomes:

$$g_{ctp}^{\text{heal}} = \frac{1}{T_{ct}} \zeta^{publ} \sum_{i=T_{ct}-L_{ct}}^{T_{ct}} \frac{y_{cip}}{(1+\delta)^i} \tag{11}$$

With y_{cip} the average income that income group p in country c can expect to receive at time i.

As documented in the existing literature, the concentration of lifetime income is substantially lower than that of current income because of significant income mobility over the life cycle (e.g., Auerbach, Kotlikoff, and Koehler, 2023). In order to estimate the value of healthcare, is it thus necessary to account for the fact that individuals belonging to a given income group are likely to move alongside the income distribution in the years to come. In the lack of good data on such income mobility patterns, I assume that lifetime income mobility follows a similar transition process as that of intergenerational mobility. I thus use the same income mobility patterns as those used for the valuation of education in the previous section. For instance, an individual at percentile 10 in Brazil can expect to receive the average income of percentile 28 at the time they benefit from greater life expectancy.

Again, I construct two alternative measures of the value of healthcare. In the future income approach, which is the benchmark specification, I use projected economic growth rates. In the current income approach, I assume that the value of an additional year of life is the current income. This is a very conservative assumption. For instance, an individual in 2022 who expects to die in 2050 instead of 2049 thanks to the healthcare system values this additional year of life as the income of 2022.

VSL Specifications As a robustness check, I also investigate directly using estimates of the value of a statistical life year (VSLY) from the literature. I do so in two steps. First, I assume that the VSLY is equal to $V_{US} = \$100,000$ in the United States in 2003 (e.g., Cutler et al., 2022). Second, I extrapolate VSLY to all other countries, years, and income groups in the data by calibrating an income elasticity ε such that: $V_{ctp} = V_{US} \times \left(\frac{y_{ctp}}{y_{US}}\right)^{\varepsilon}$, with y_{US} U.S. GDP per capita in 2003. $\varepsilon = 1$ implies adjusting VSLY proportionally to income. $\varepsilon < 1$ implies that VSLY is concave in income, while $\varepsilon > 1$ implies that VSLY is convex in income. I consider alternative values of 0.5, 1, and 1.5, in line with ranges typically found in the literature (e.g., Hammitt and Robinson, 2011).

B. Additional Figures and Tables

B.1. Additional Main Results

Figure A1 – Validation of Methodology: Trends in Regional Healthcare Inequality.
Coefficient of Variation in Subnational Healthcare Quality Index, 1990-2016

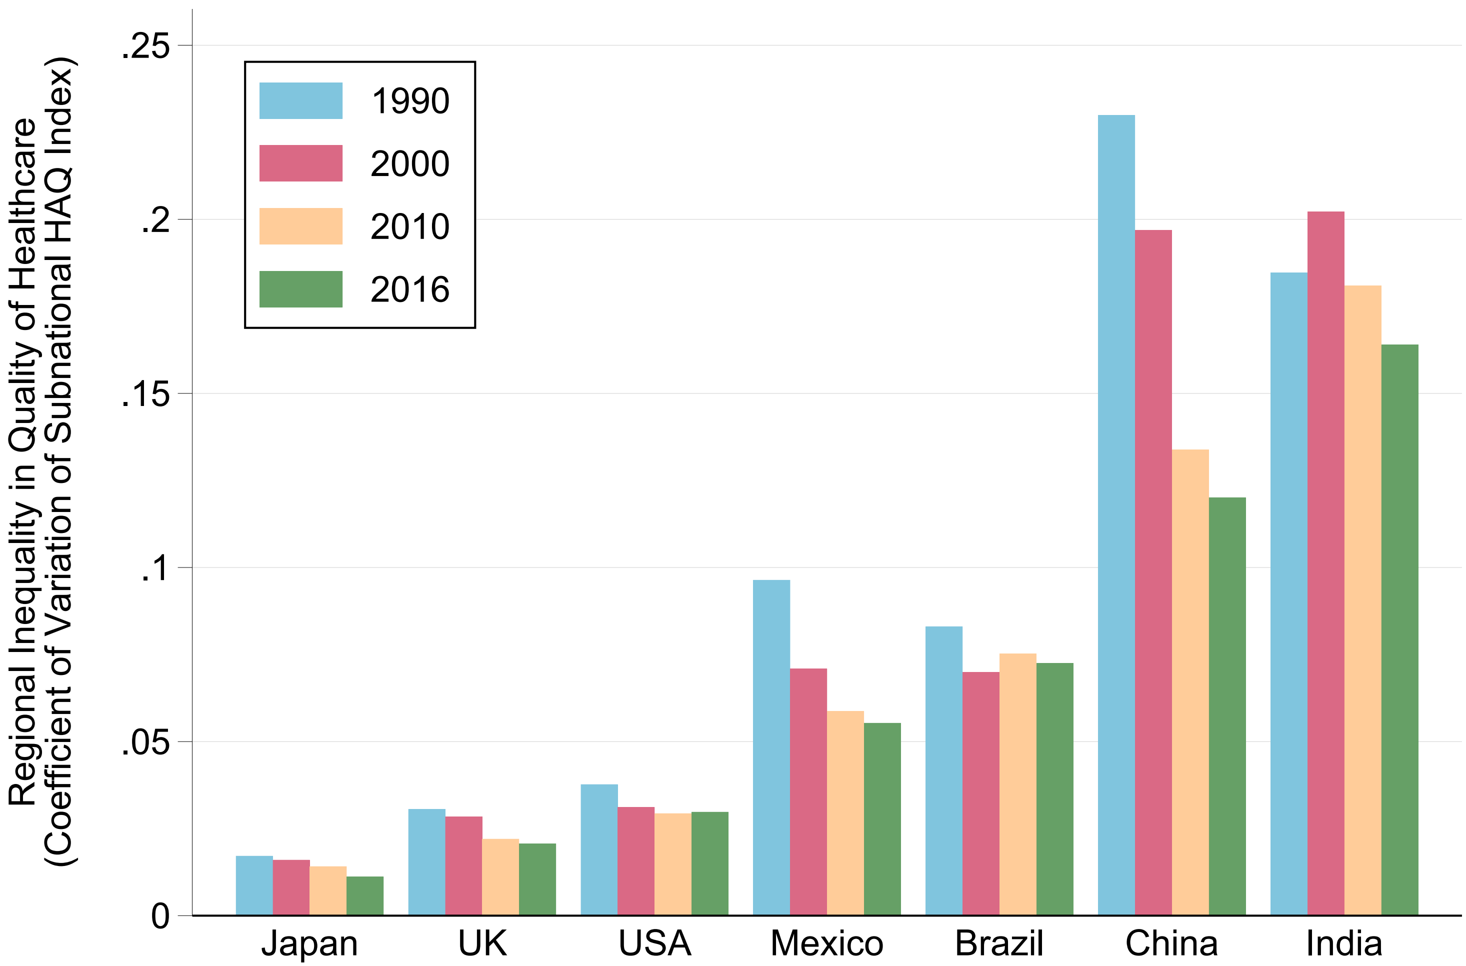

Notes. The figure plots the evolution of regional inequality in the quality of healthcare in selected countries, measured by the coefficient of variation of the subnational HAQ index available from the Global Burden of Disease study (GBD, 2022).

Figure A2 – Global Real Public Expenditure Per Capita (2021 PPP USD)

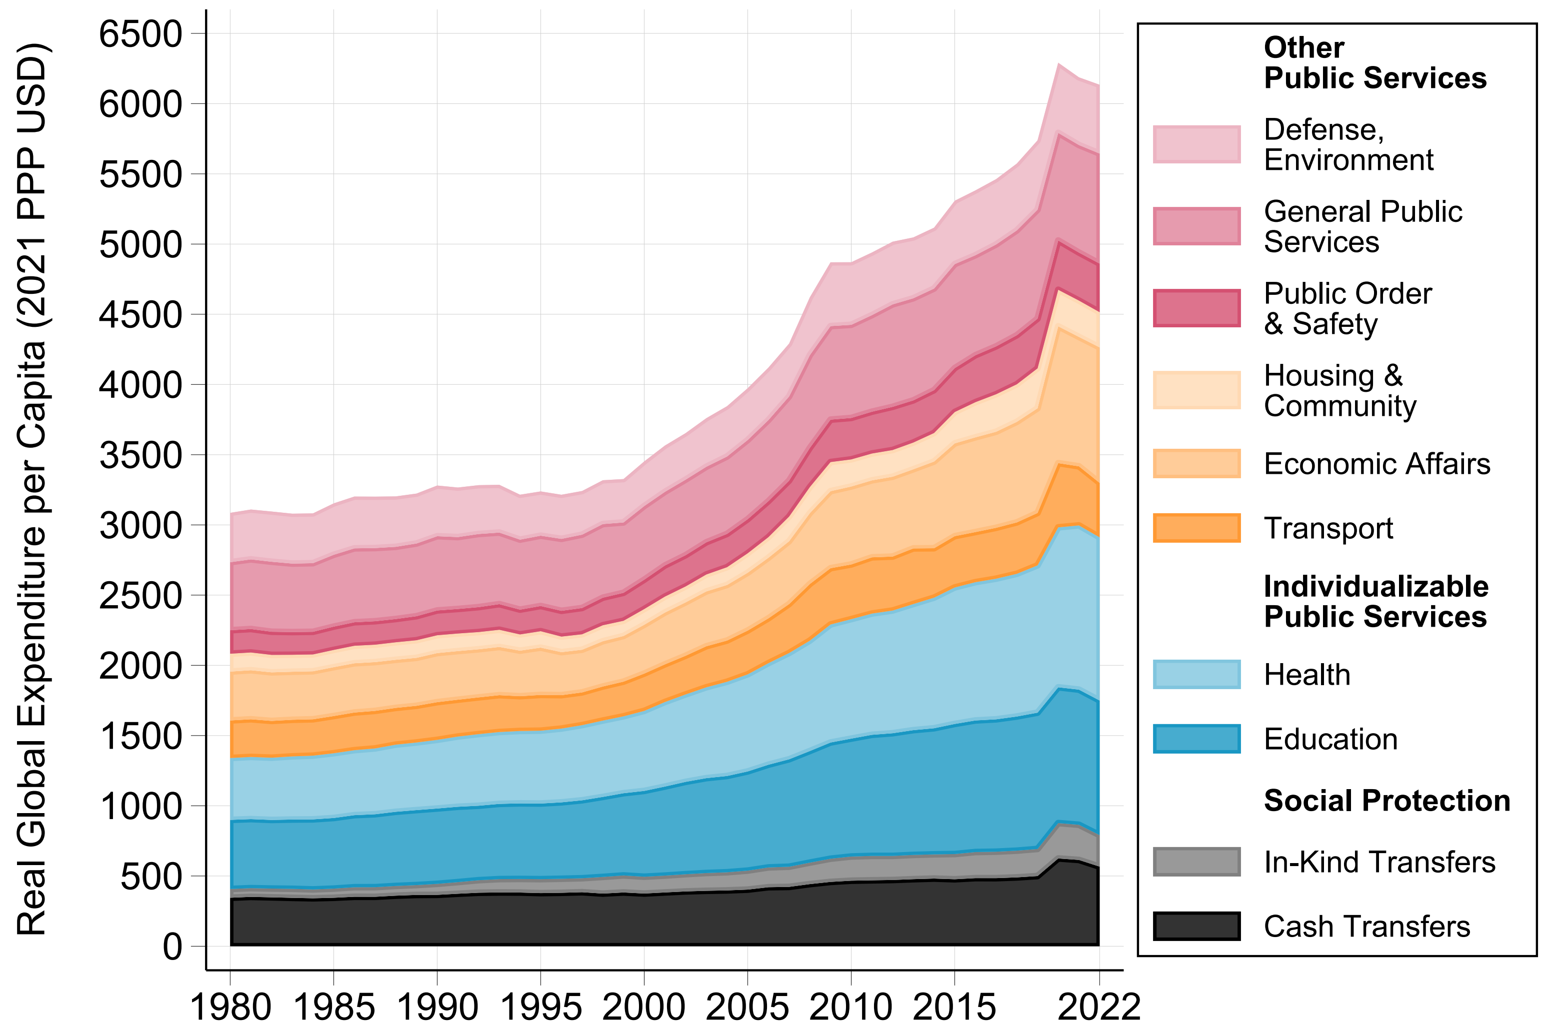

Notes. The figure represents the evolution of general government expenditure per capita by function, expressed in 2021 PPP US dollars, in the world as a whole. Average expenditure per capita grew from \$3000 to \$6000 from 1980 to 2022. Social protection: cash transfers and in-kind social assistance transfers such as food stamps. Transport: expenditure on public transport and transport infrastructure. Economic affairs: expenditure on economic affairs other than transport, such as subsidies to agriculture, energy, manufacturing, and recreation and culture. Housing and community: expenditure on public housing and community services such as water supply and electricity. Public order and safety: expenditure on police services, law courts, and prisons. Author's computations based on historical budget data collected in this paper.

Figure A3 – Government Transfers Received by the World's Poorest 20% (Real 2021 PPP USD)

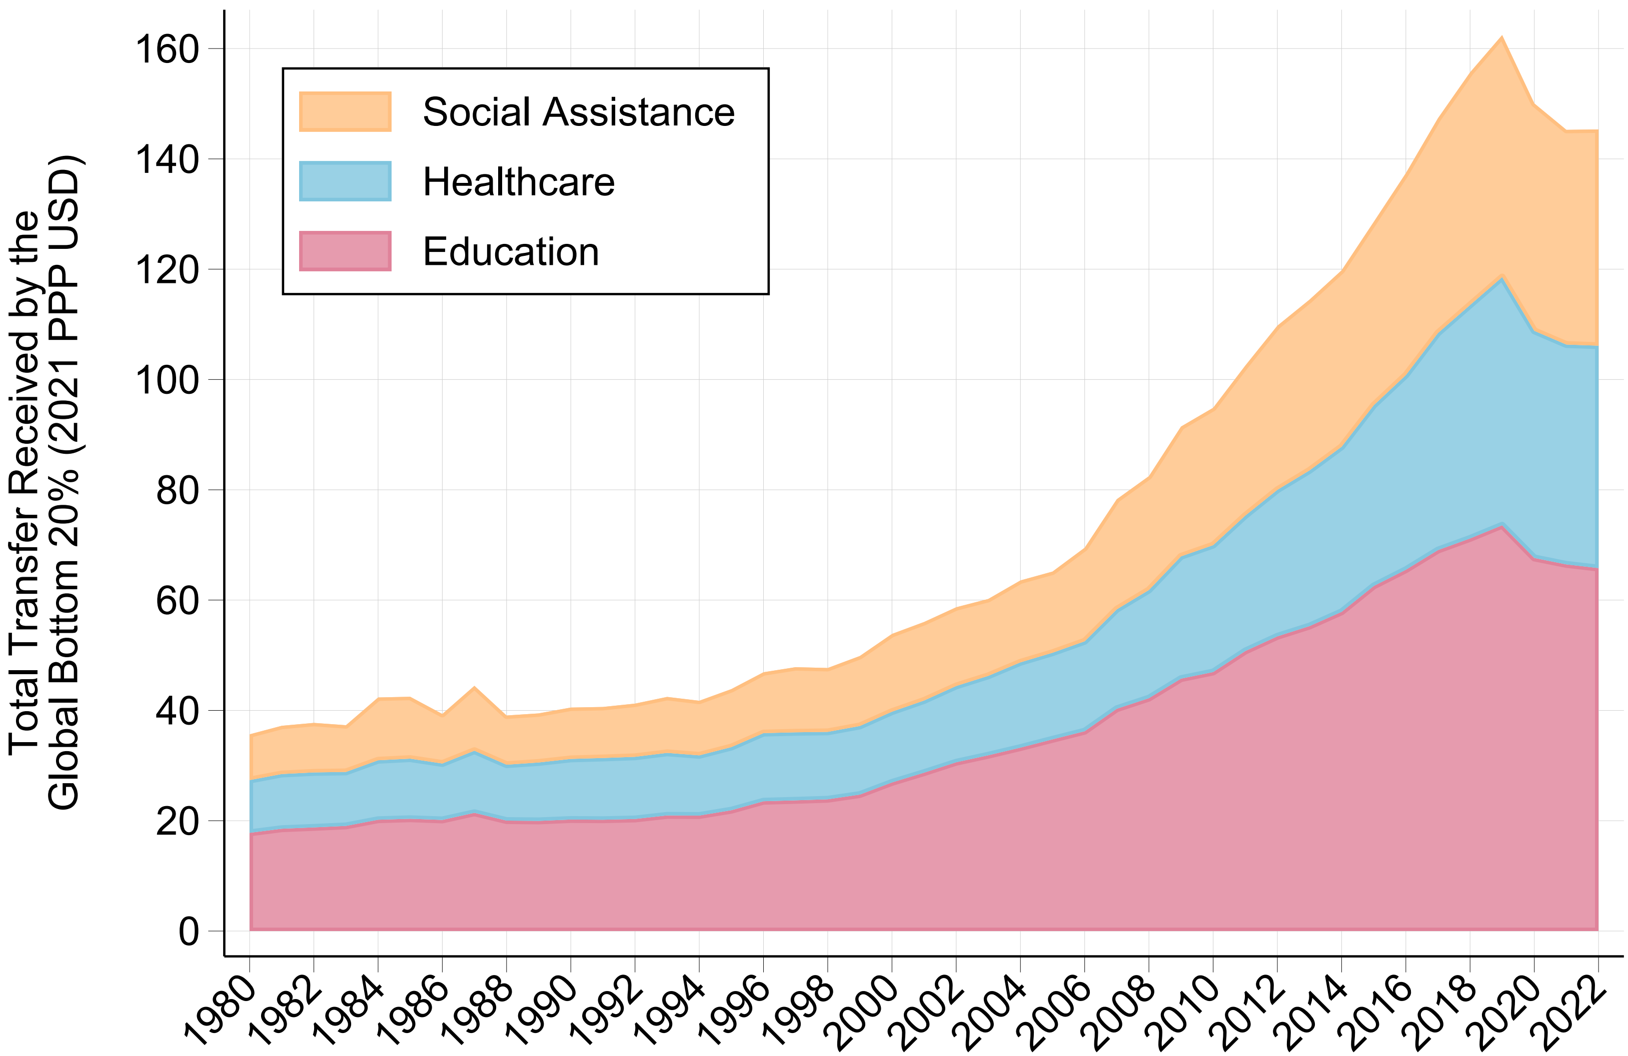

Notes. The figure plots the level and composition of government transfers received by the world's poorest 20%, expressed in real 2021 PPP US dollars, from 1980 to 2022.

Figure A4 – Government Redistribution and Global Poverty Reduction:
Global Poverty Headcount Ratio at $2.15 per Day, 1980-2022

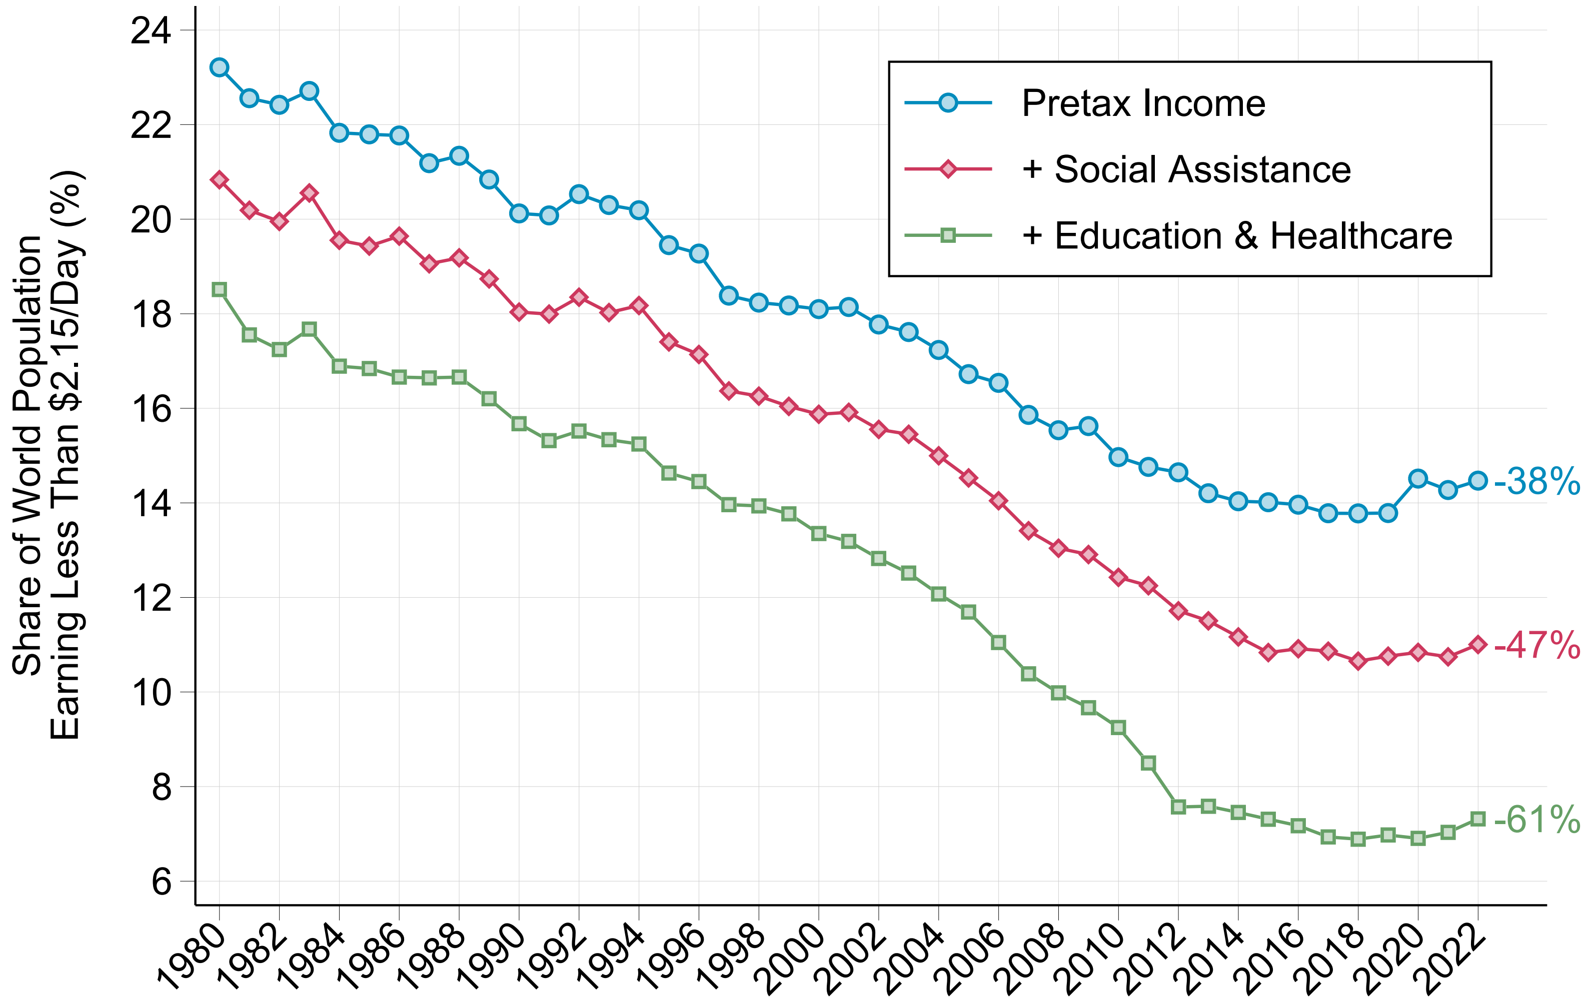

Notes. The figure plots the evolution of the poverty headcount ratio at $2.15 per day (2017 PPP USD) in the world as a whole, before and after adding social assistance transfers and the consumption of public services to individual incomes. Global poverty declined by 38% in terms of pretax income, compared to 47% after accounting for social assistance transfers, and 61% after adding social assistance, education, and healthcare transfers to individual incomes. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare.

Figure A5 – Government Redistribution and Global Inequality:
Global Top 10% to Bottom 50% Average Income Ratio

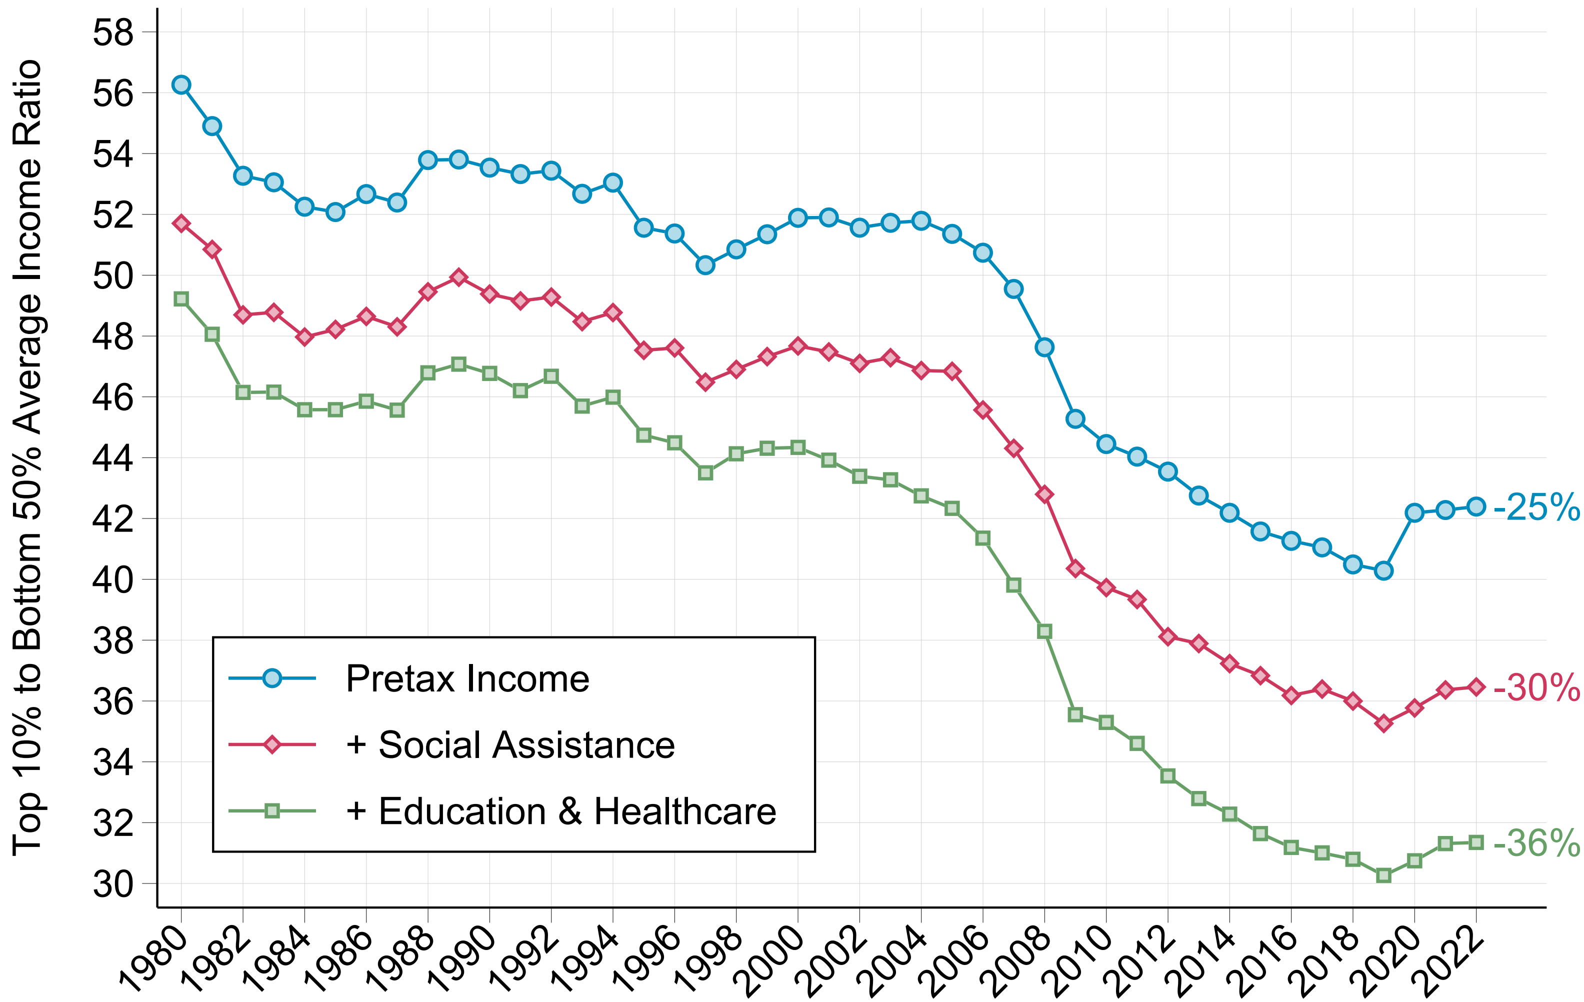

Notes. The figure plots the evolution of global income inequality, measured as the ratio of the average income of the top 10% to that of the bottom 50% in the world as a whole, before and after adding social assistance transfers and the consumption of public services to individual incomes. Worldwide inequality declined by 25% in terms of pretax income, compared to 36% after accounting for all government transfers. Social assistance includes cash transfers and in-kind social benefits such as food stamps. Public education and healthcare are valued at cost of provision, that is, as total general government expenditure on education and healthcare.

Figure A6 – Gini Index of Global Income Inequality

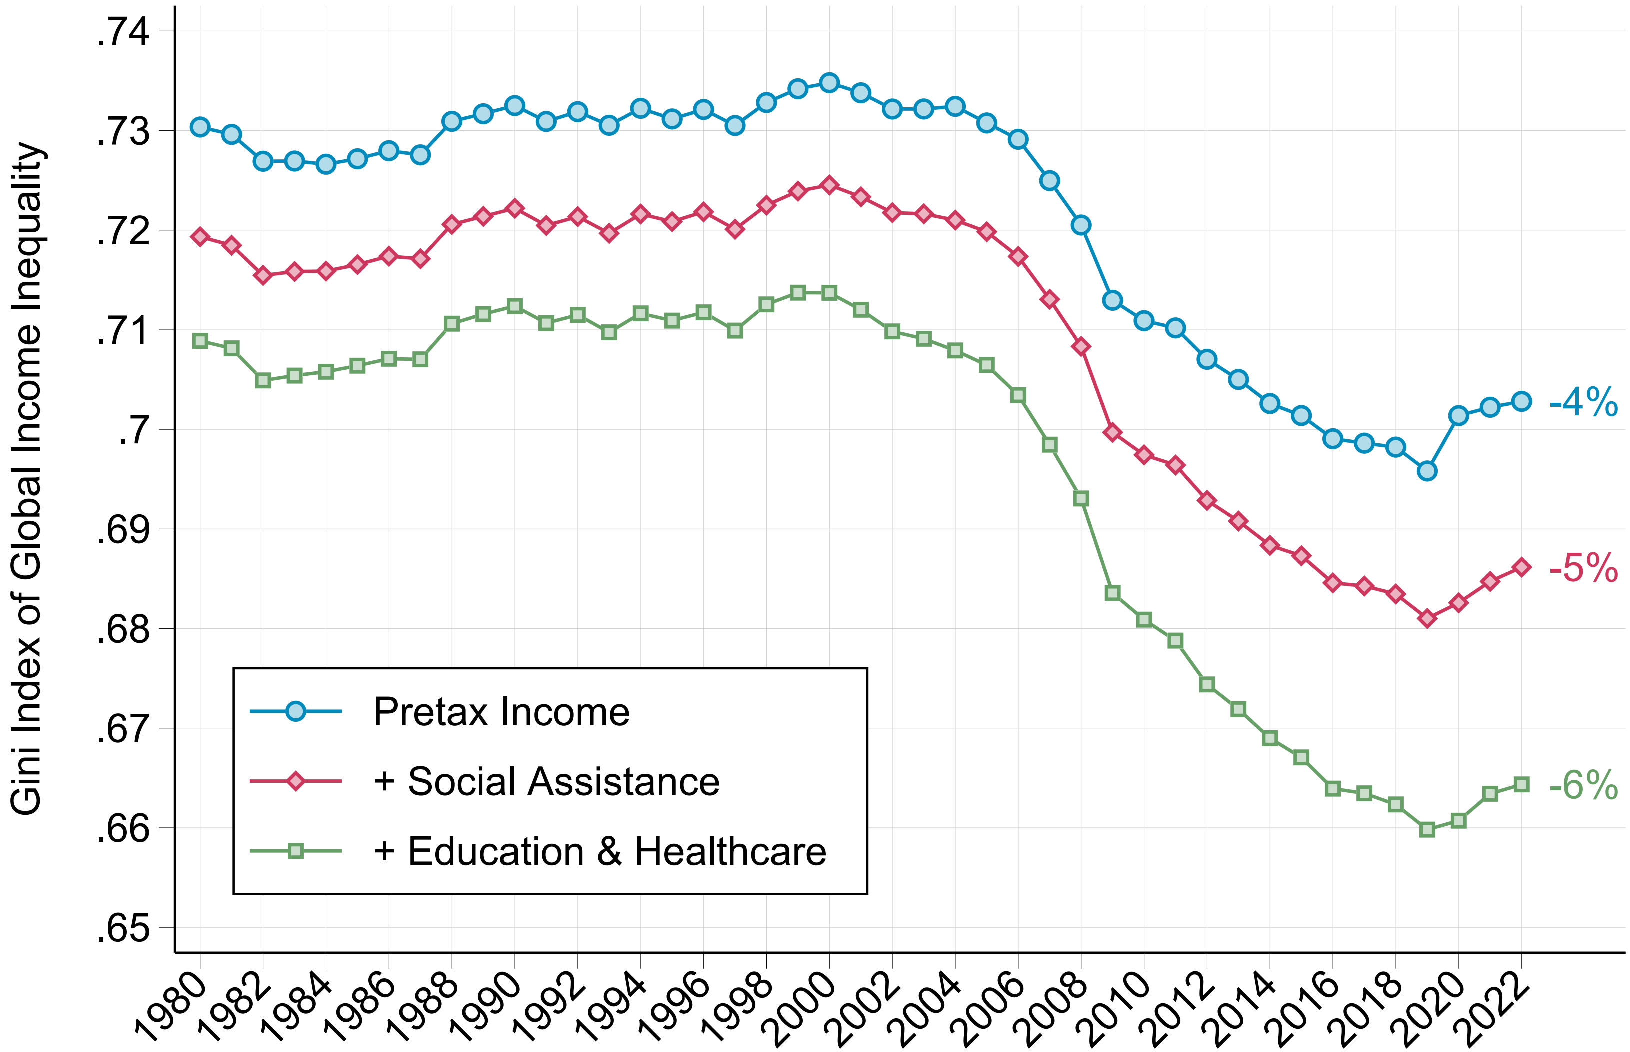

Notes. The figure plots the evolution of the Gini index of global income inequality, before and after adding social assistance transfers and the consumption of public services to individual incomes.

Figure A7 – Theil Index of Global Income Inequality

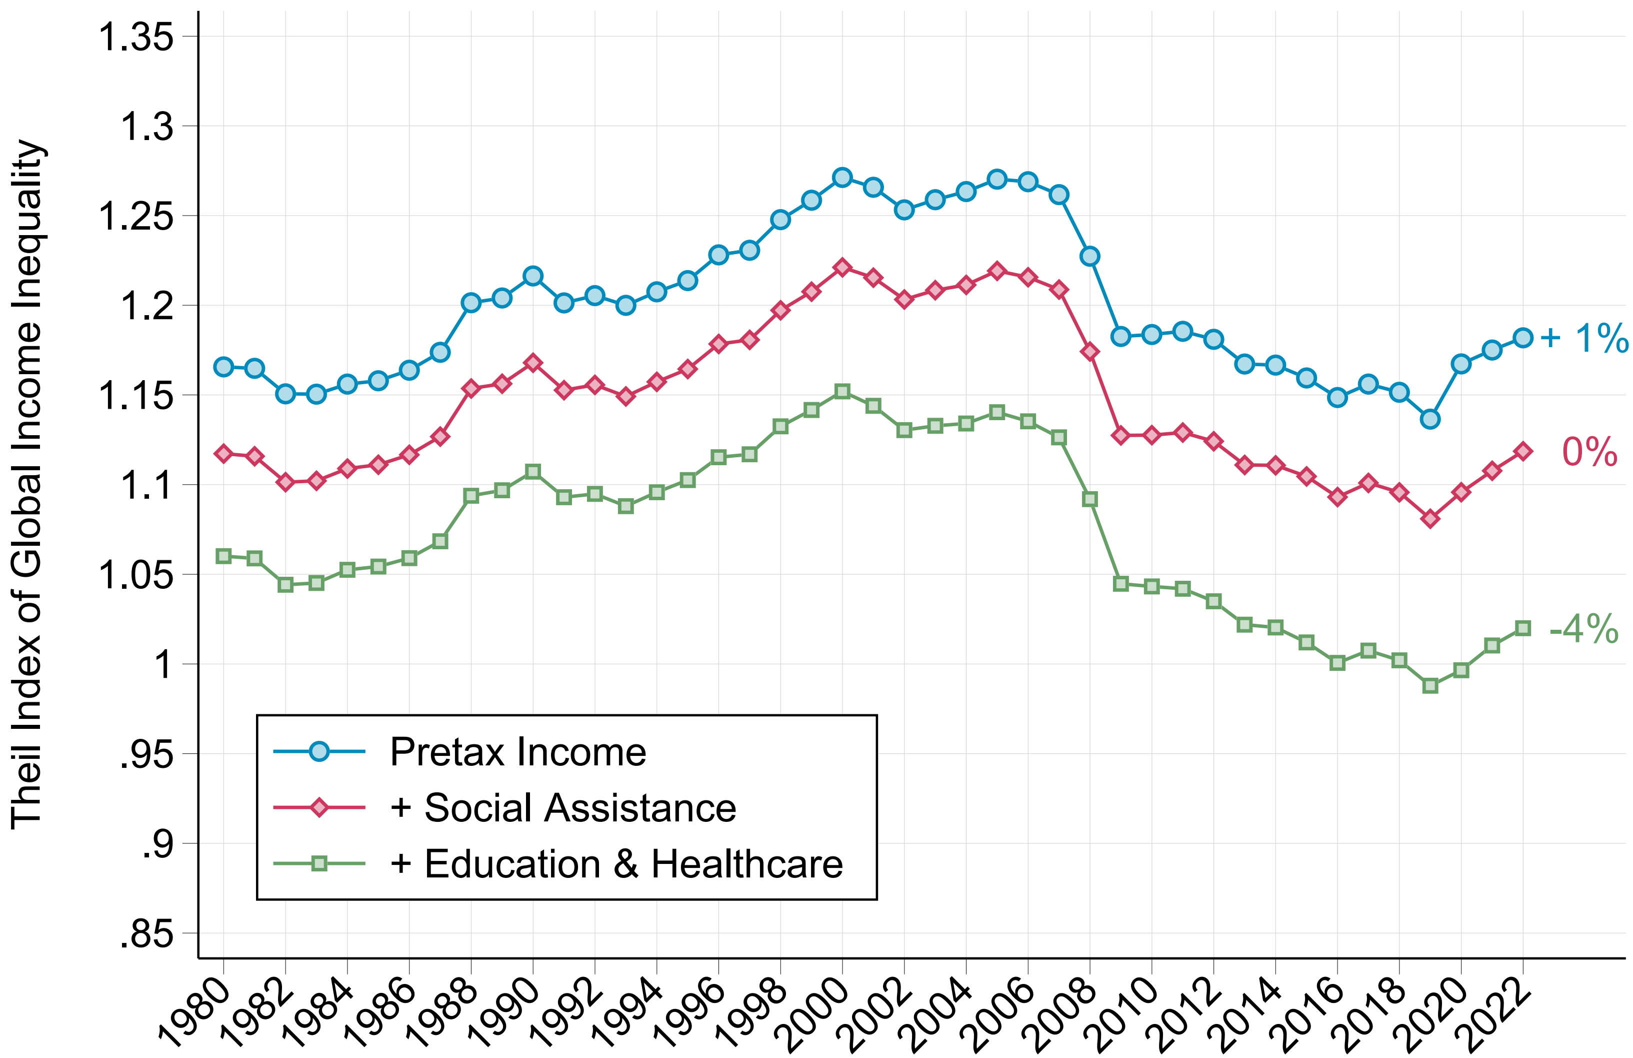

Notes. The figure plots the evolution of the Theil index of global income inequality, before and after adding social assistance transfers and the consumption of public services to individual incomes.

Table A1 – Government Redistribution and Global Bottom 20% Growth: Sensitivity to Different Specifications and Geographical Restrictions

	Global Bottom 20% Average Income (2021 PPP USD)		
	1980	2022	2022-1980
All Countries			
Pretax Income	325	580	78%
Posttax Income: Benchmark	508	1392	174%
Posttax Income: Other Public Services Lump Sum	688	1870	172%
Excluding China			
Pretax Income	319	496	55%
Posttax Income	531	1156	118%
Excluding India			
Pretax Income	363	539	48%
Posttax Income	593	1365	130%
Excluding China & India			
Pretax Income	360	436	21%
Posttax Income	695	1061	53%

Notes. The table reports how results on the incidence of government redistribution on real income growth of the world's poorest 20% vary depending on assumptions regarding the progressivity of public services and geographical restrictions. Posttax income equals pretax income, plus all transfers (social assistance, education, healthcare, and other public services), minus all taxes (direct and indirect taxes). Benchmark Posttax Income: all government expenditure other than education, health, and social assistance allocated proportionally to disposable income. Other Public Services Lump Sum: all government expenditure other than education, health, and social assistance allocated as a lump sum.

Table A2 – Government Redistribution and Global Poverty Reduction: Results Using World Bank Income Distribution Data

	1980	2022	2022-1980
Global Bottom 20% Transfer (% Global Income)			
Social Assistance	0.6%	2.6%	+303%
Education	1.1%	3.7%	+236%
Healthcare	0.6%	1.8%	+198%
All Transfers	1.7%	5.5%	+222%
Bottom 20% Average Income ($2017 PPP/Day)			
Pretax Income	0.5	1.5	+184%
+ Social Assistance	0.7	2.2	+219%
+ Education & Healthcare	0.8	2.9	+255%
Posttax Income	0.7	2.8	+269%
Global Poverty Headcount Ratio ($2.15/Day)			
Pretax Income	46.1%	14.3%	-69%
+ Social Assistance	42.4%	9.0%	-79%
+ Education & Healthcare	38.9%	5.2%	-87%
Posttax Income	41.2%	5.5%	-87%

Notes. The table reports results on transfers received by the world's poorest 20%, the average income of the global bottom 20%, and the global poverty rate from 1980 to 2022. The world distribution of income is reconstructed using data from the World Bank's Poverty and Inequality Platform. Posttax income equals pretax income, plus all transfers (social assistance, education, healthcare, and other public services), minus all taxes (direct and indirect taxes).

B.2. Accounting for Public Sector Productivity

Figure A8 – Expected Human Capital from Age 5 Onward by Country Income Group, 1980-2022

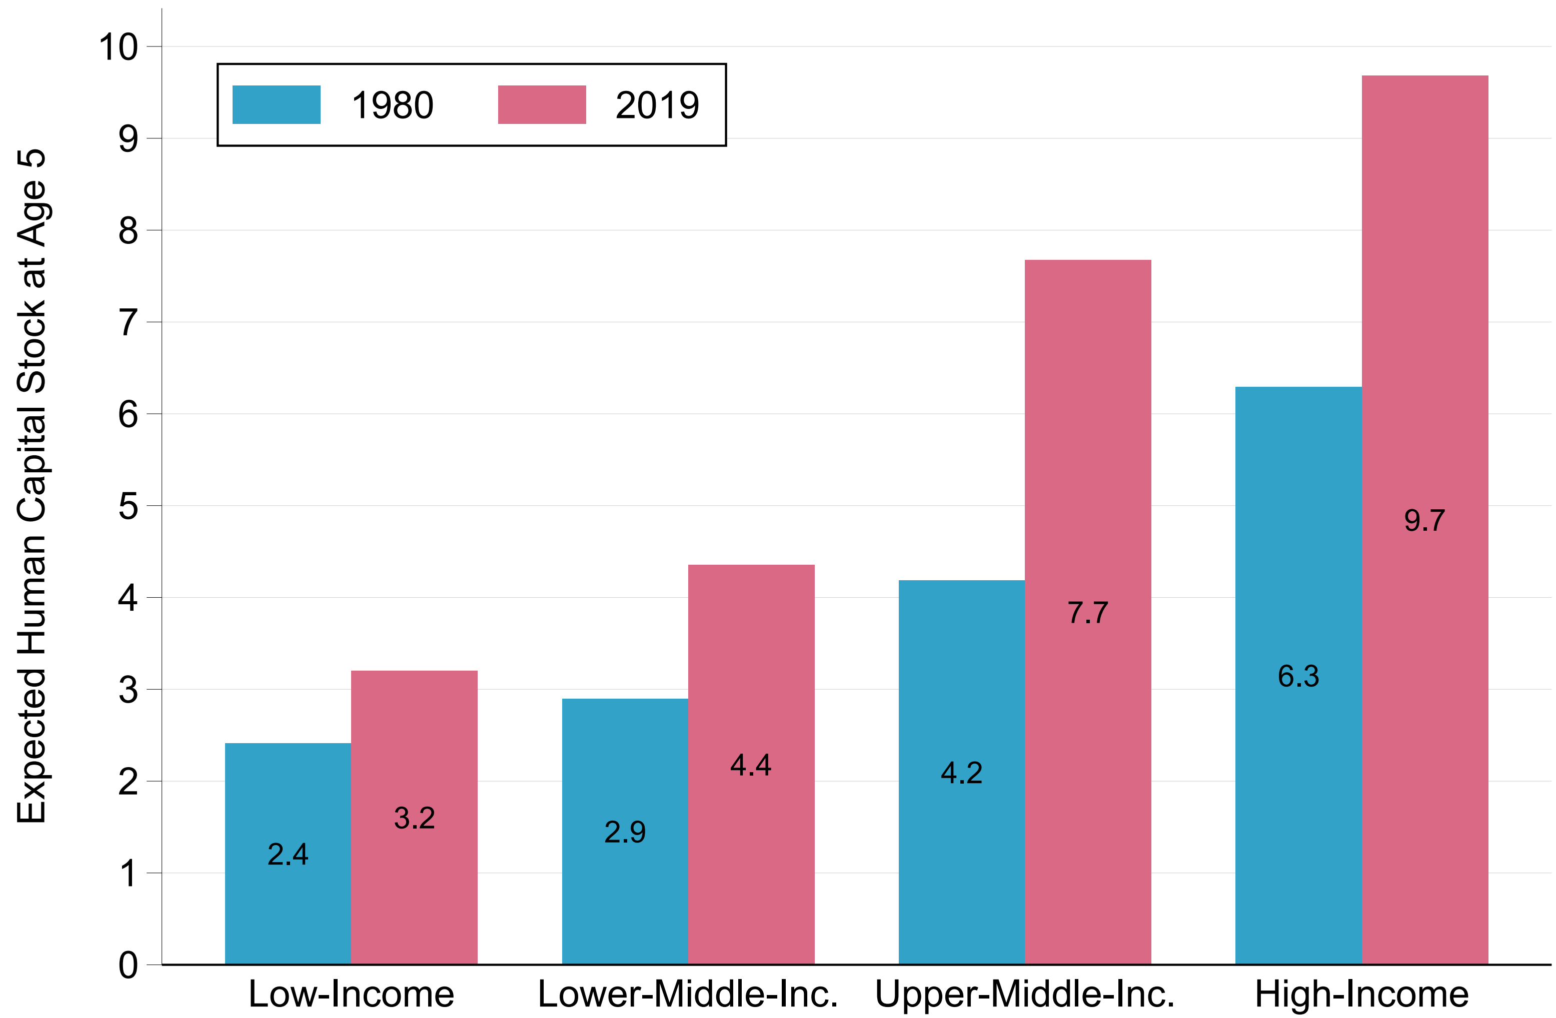

Notes. The figure plots the evolution of expected human capital at age 5 by country income group. Expected human capital is calculated as $Y^{\text{education}} = \exp(r_S S + r_Q Q)$, with S expected years of schooling at age 5, r_S the return to a year of schooling, Q standardized education quality, and r_Q the return to standardized education quality. Population-weighted averages across all countries in each group.

Figure A9 – Healthcare Access and Quality Index by Country Income Group, 1990-2016

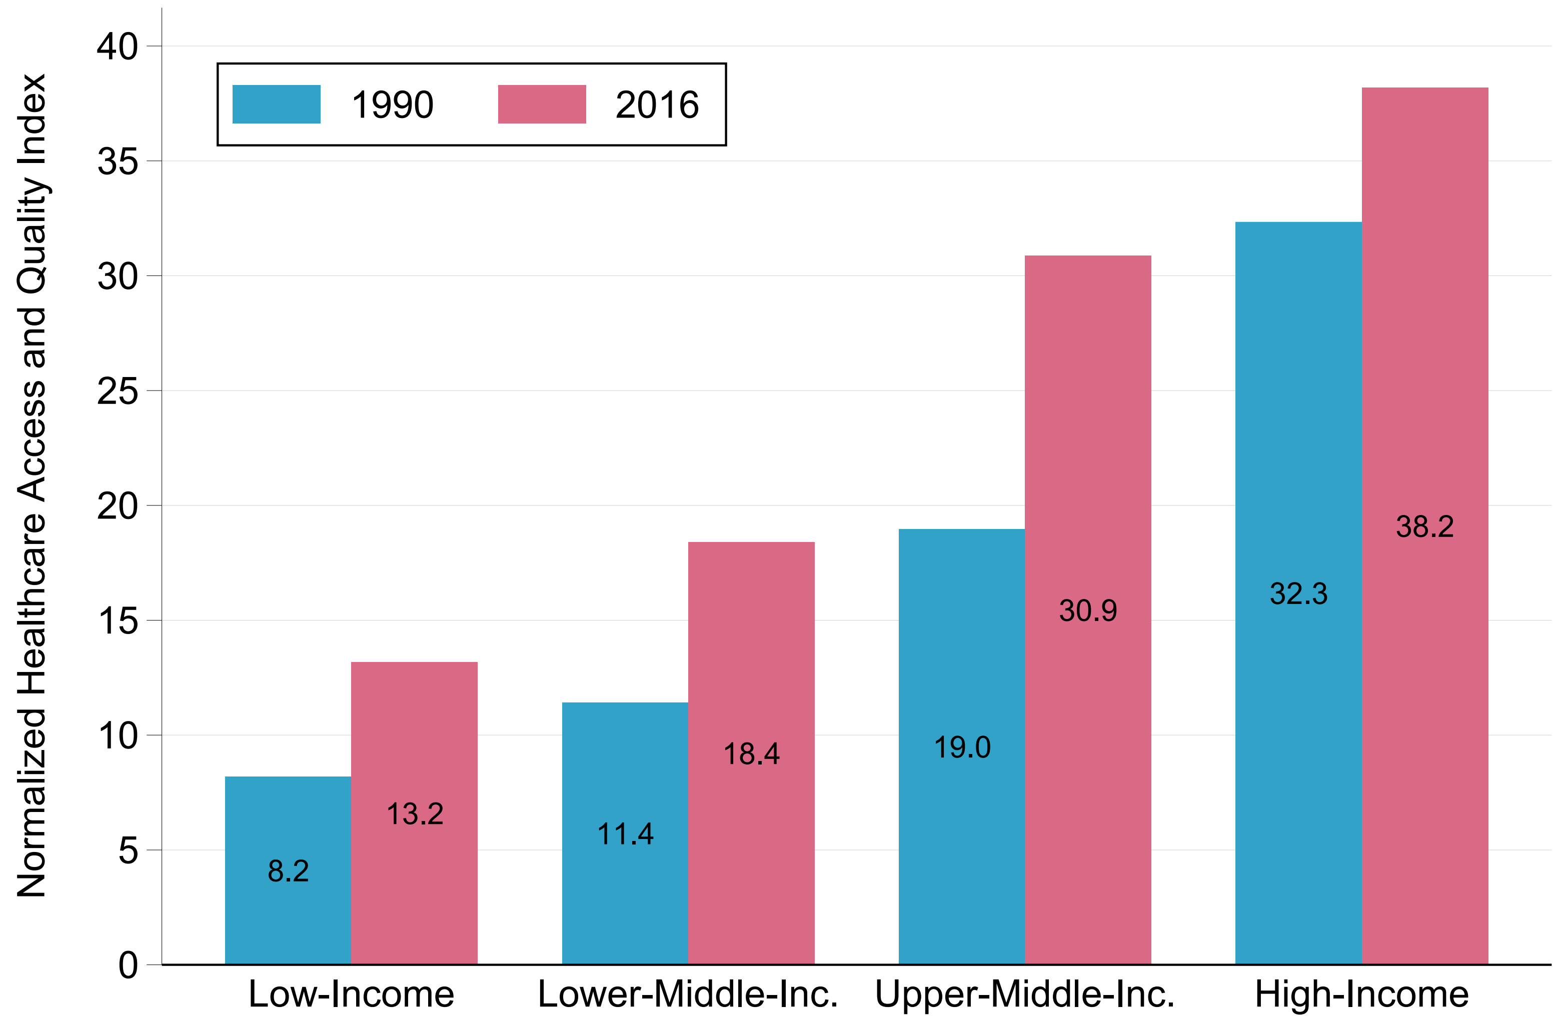

Notes. The figure plots the evolution of the healthcare access and quality index by country income group. The index is taken from GBD (2022) and reexpressed in units of life expectancy gains enabled by the healthcare system. Population-weighted averages across all countries in each group.

Figure A10 – Total Public Education and Healthcare Transfers Received
by the Global Bottom 20%: Before Versus After Productivity Adjustment

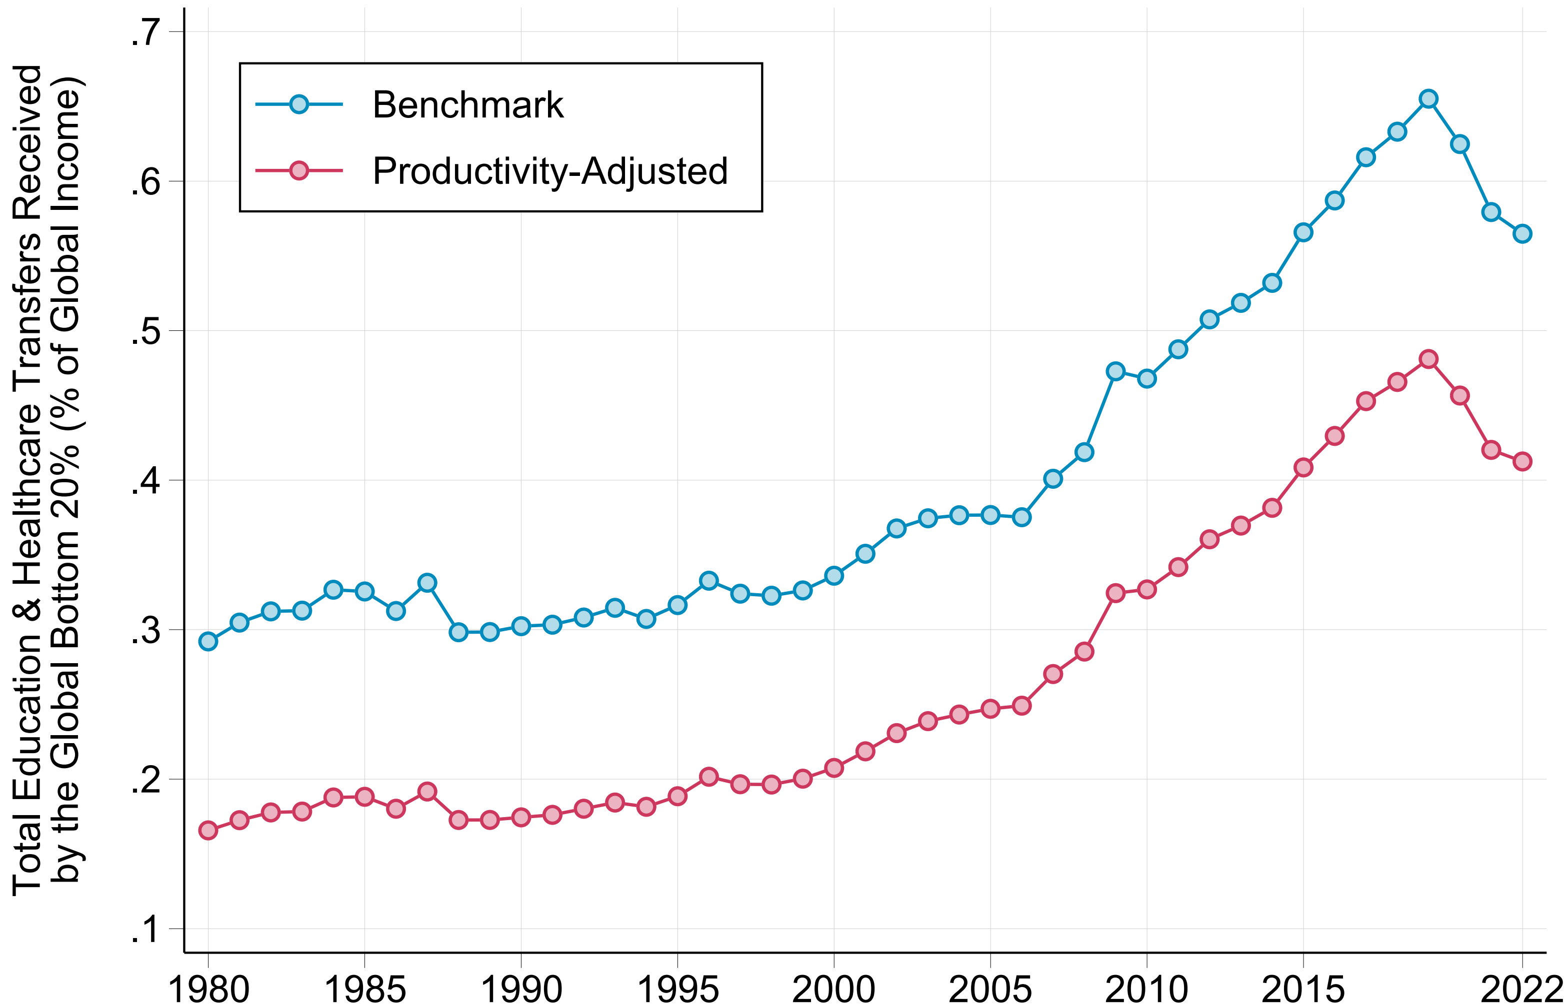

Notes. The figure plots the level and composition of government transfers received by the world's poorest 20%, expressed as a share of global income, before and after adjusting public services for public sector productivity.

Figure A11 – Real Average Income of the Global Bottom 20%: Before versus After Productivity Adjustment

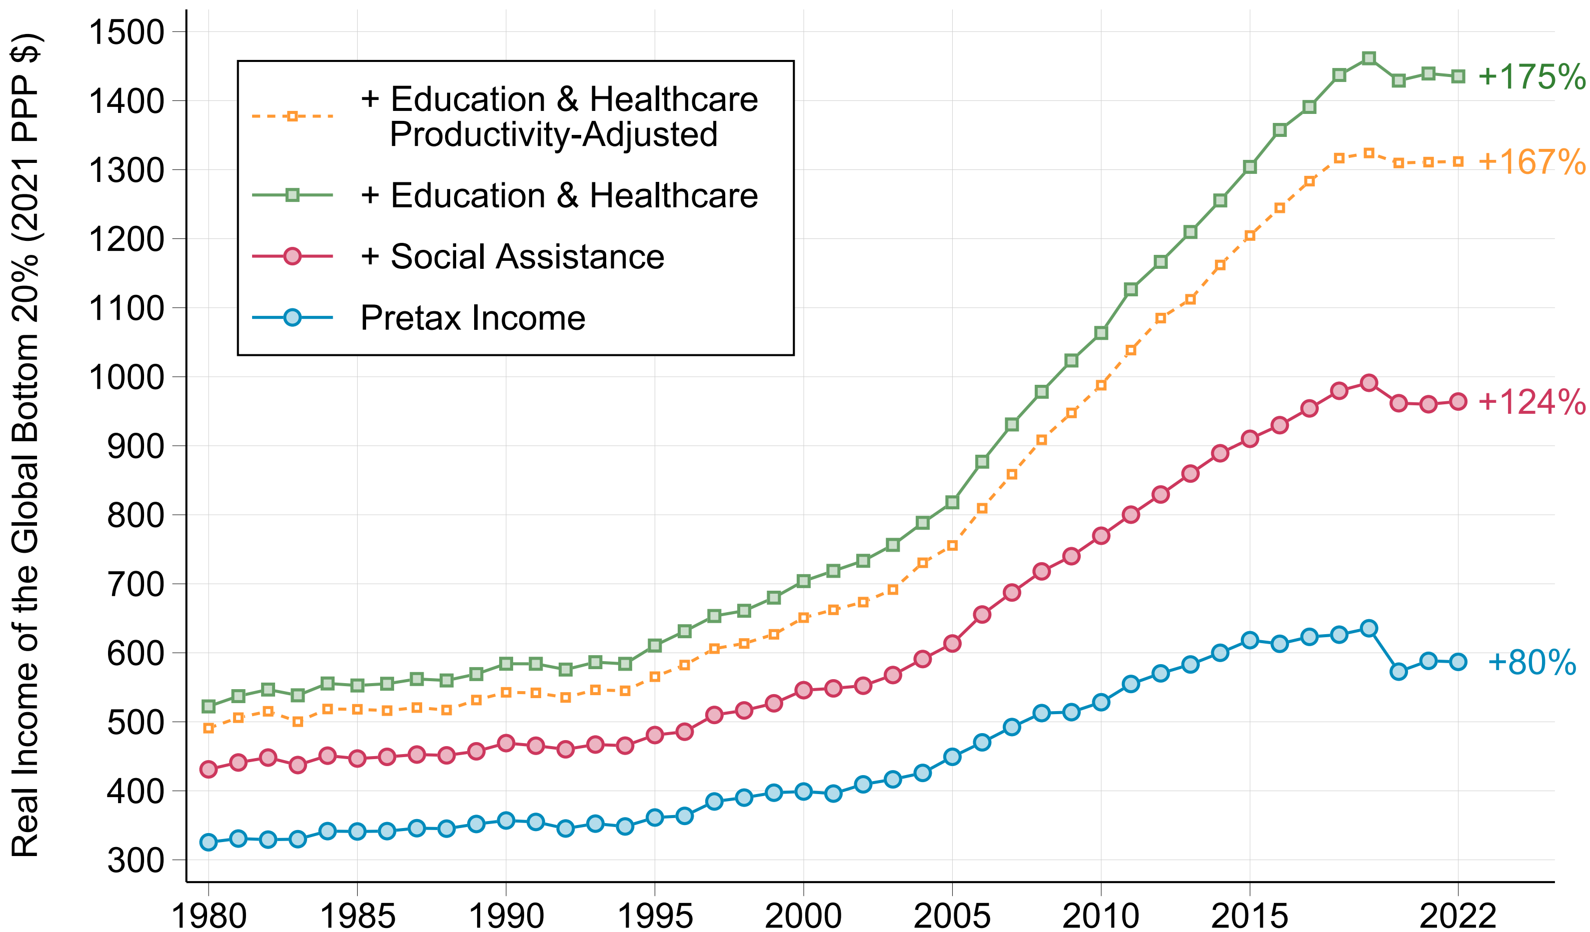

Notes. The figure plots the evolution of the global bottom 20% real average income from 1980 to 2022, before and after adding social assistance transfers and the consumption of public education and healthcare to individual incomes. Social assistance includes cash transfers and in-kind social benefits such as food stamps. The solid line values public education and healthcare at cost of provision, that is, as total general government expenditure on education and healthcare. The dashed line adjusts estimates of the consumption of public services for differences in public sector productivity.

Figure A12 – Real Income Growth Rate by Global Income Percentile, 1980-2022: Before versus After Productivity Adjustment

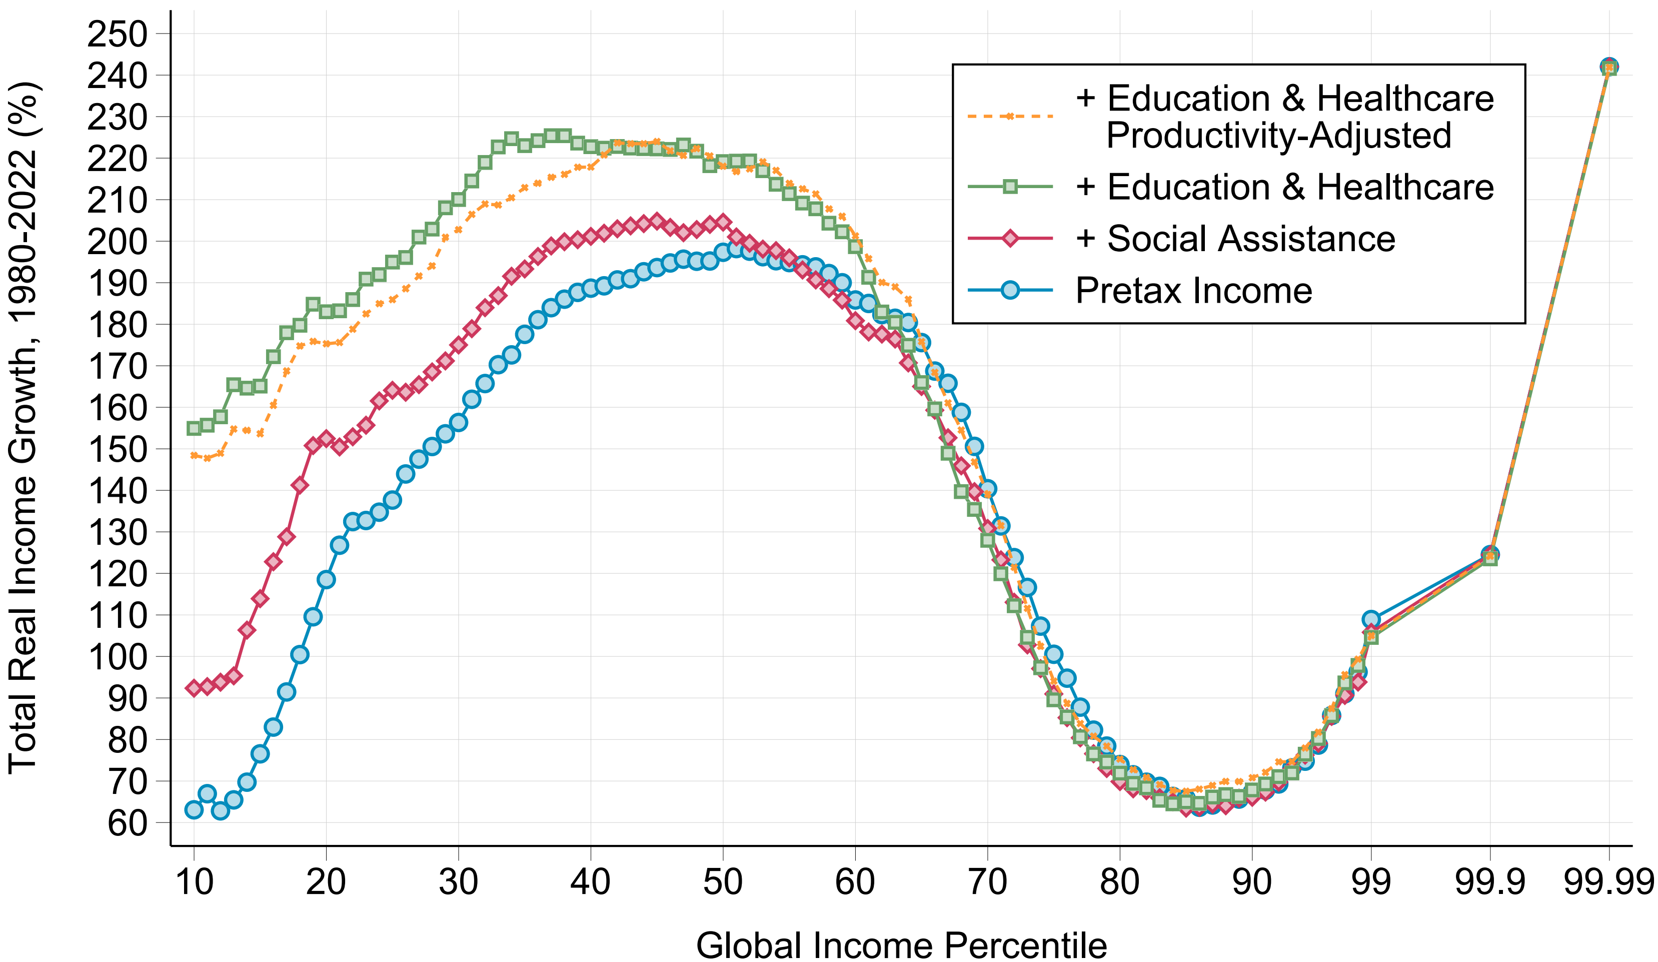

Notes. The figure plots total real income growth by global income percentile from 1980 to 2022, before and after adding social assistance transfers and the consumption of public services to individual incomes. The dashed line adjusts estimates of the consumption of public services for differences in public sector productivity.

Table A3 – Government Productivity by Country Income Group (Efficient Frontier = 1)

	Education		Healthcare	
	1980	2022	1980	2022
Low-Income	0.51	0.63	0.29	0.59
Lower-Middle-Income	0.61	0.62	0.42	0.64
Upper-Middle-Income	0.76	0.75	0.68	0.89
High-Income	0.58	0.72	0.84	0.93
All Countries	0.65	0.68	0.60	0.76

Notes. The table reports average public sector productivity (Θ^j) in education and healthcare by country income group in 1980 and 2022. Population-weighted averages of efficiency scores in each country. Country-years with a score of 1 are those at the efficient frontier.

Table A4 – Correlates of Public Sector Productivity

	Education	Healthcare	N
Chong et al. (2014) Mail Efficiency	0.27***	0.61***	159
Government Effectiveness	0.45***	0.66***	177
Control of Corruption	0.44***	0.58***	177
Absence of Corruption	0.40***	0.54***	160
Wastefulness of Government Spending	0.18**	0.20**	149
Irregular Payments and Bribes	0.45***	0.61***	150
Favoritism in Government Decisions	0.39***	0.34***	151
Transparency of Policymaking	0.38***	0.43***	150
GDP per capita	0.37***	0.61***	177

Notes. The table reports raw pairwise correlations between the two measures of public sector productivity developed in this paper and other qualitative indicators of government productivity from the available literature. Correlations are computed over all countries with available data for each pair of indicators, for the last year available. Chong et al. (2014) efficiency: average number of days to get the letter back from a given country. GDP per capita data come from the World Inequality Database. Data on other indicators come from the World Bank. * p<0.10, ** p<0.05, *** p<0.01.

Table A5 – Government Redistribution Over the Course of Development: Before Versus After Adjusting for Public Sector Productivity

	Expenditure (% NNI) G	Share of Transfer Received (%) (γ, Bottom 50%)	Net Transfer Received (% NNI) (g, Bottom 50%)	Adjusted for Productivity (g, Bottom 50%)
Country Income Group				
Low-Income	7.0%	49.2%	3.4%	2.4%
Lower-Middle-Income	7.9%	50.3%	4.0%	2.8%
Upper-Middle-Income	11.1%	52.4%	5.9%	5.1%
High-Income	23.4%	59.1%	14.0%	12.7%
World Region				
Sub-Saharan Africa	7.6%	48.1%	3.7%	2.6%
Middle East and Northern Africa	9.7%	51.8%	5.0%	3.9%
China	8.9%	50.5%	4.5%	4.1%
India	9.2%	49.9%	4.6%	3.1%
Other Asia / Oceania	9.0%	53.2%	4.9%	4.1%
Latin America	15.4%	56.2%	8.8%	7.2%
US / Canada / Western Europe	22.2%	58.4%	13.2%	11.9%

Notes. The table reports statistics on dimensions of redistribution by country income group and world region. All figures focus on total redistribution in the form of social assistance, education, and healthcare. The fourth column reports the total transfer received by the bottom 50%, expressed as a share of national income. The last column adjusts estimates for differences in productivity across countries.

B.3. Welfare Value of Public Services

Figure A13 – Welfare Value of Public Services by Country Income Group

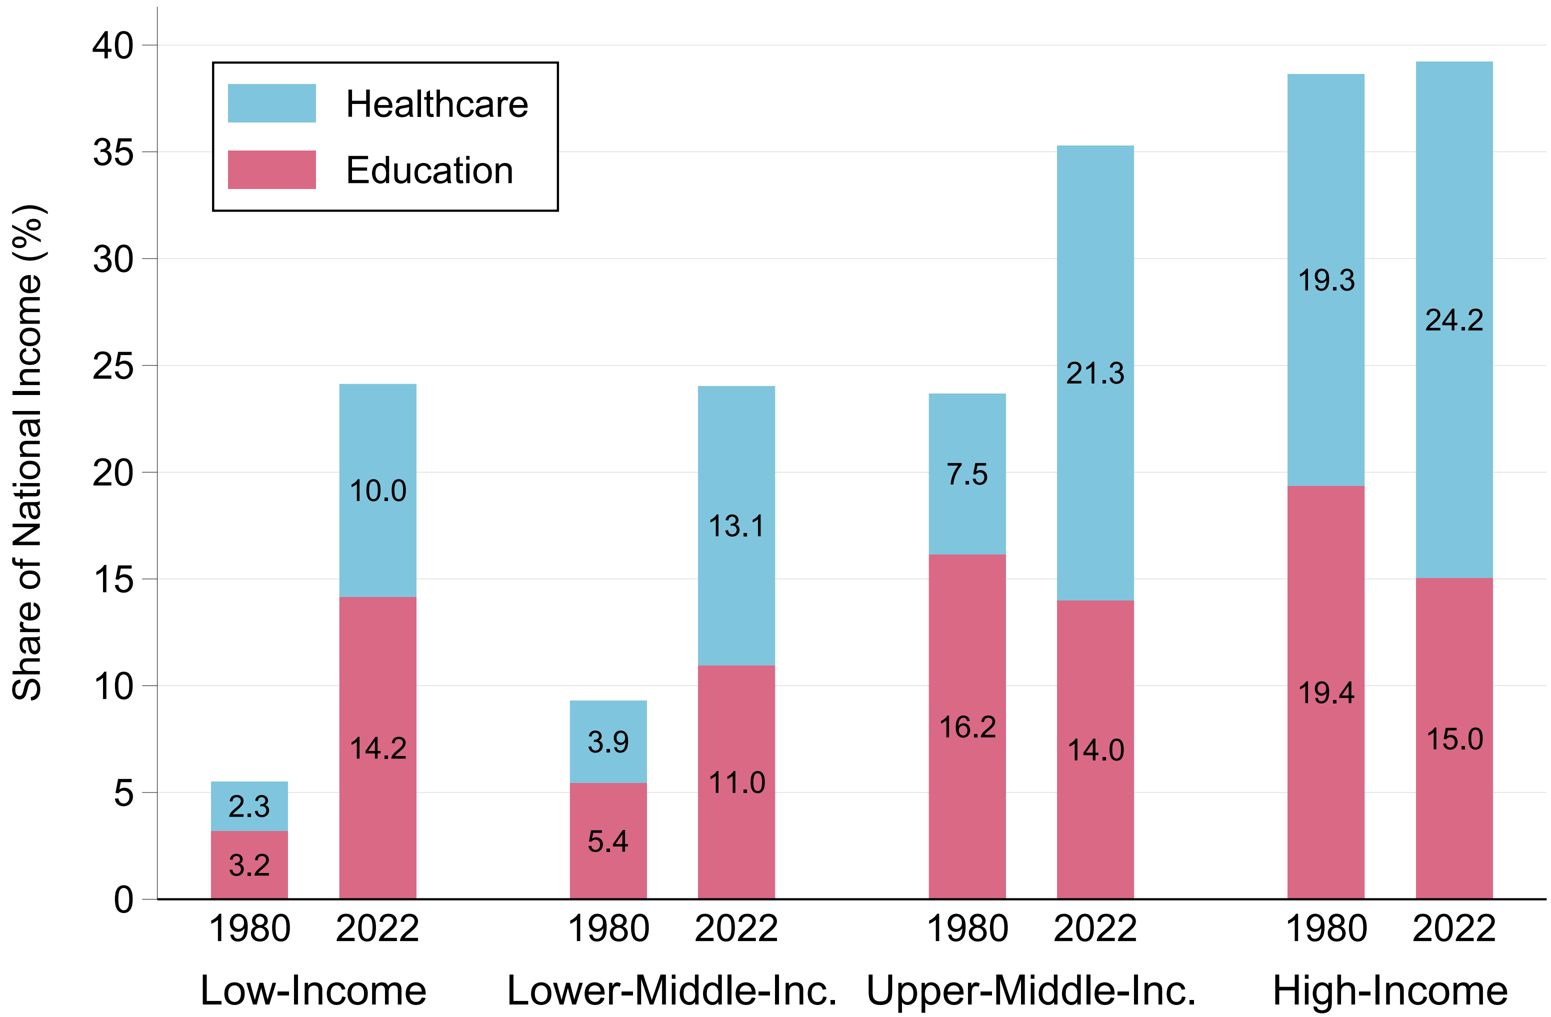

Notes. The figure plots the welfare value of public education and healthcare for the population as a whole, expressed as a share of national income, by country income group. The welfare value of education and healthcare is estimated using the future income approach, that is, accounting for future economic growth. Population-weighted averages across all countries in each group.

Figure A14 – Welfare Value of Public Services by Country Income Group, Bottom 50%

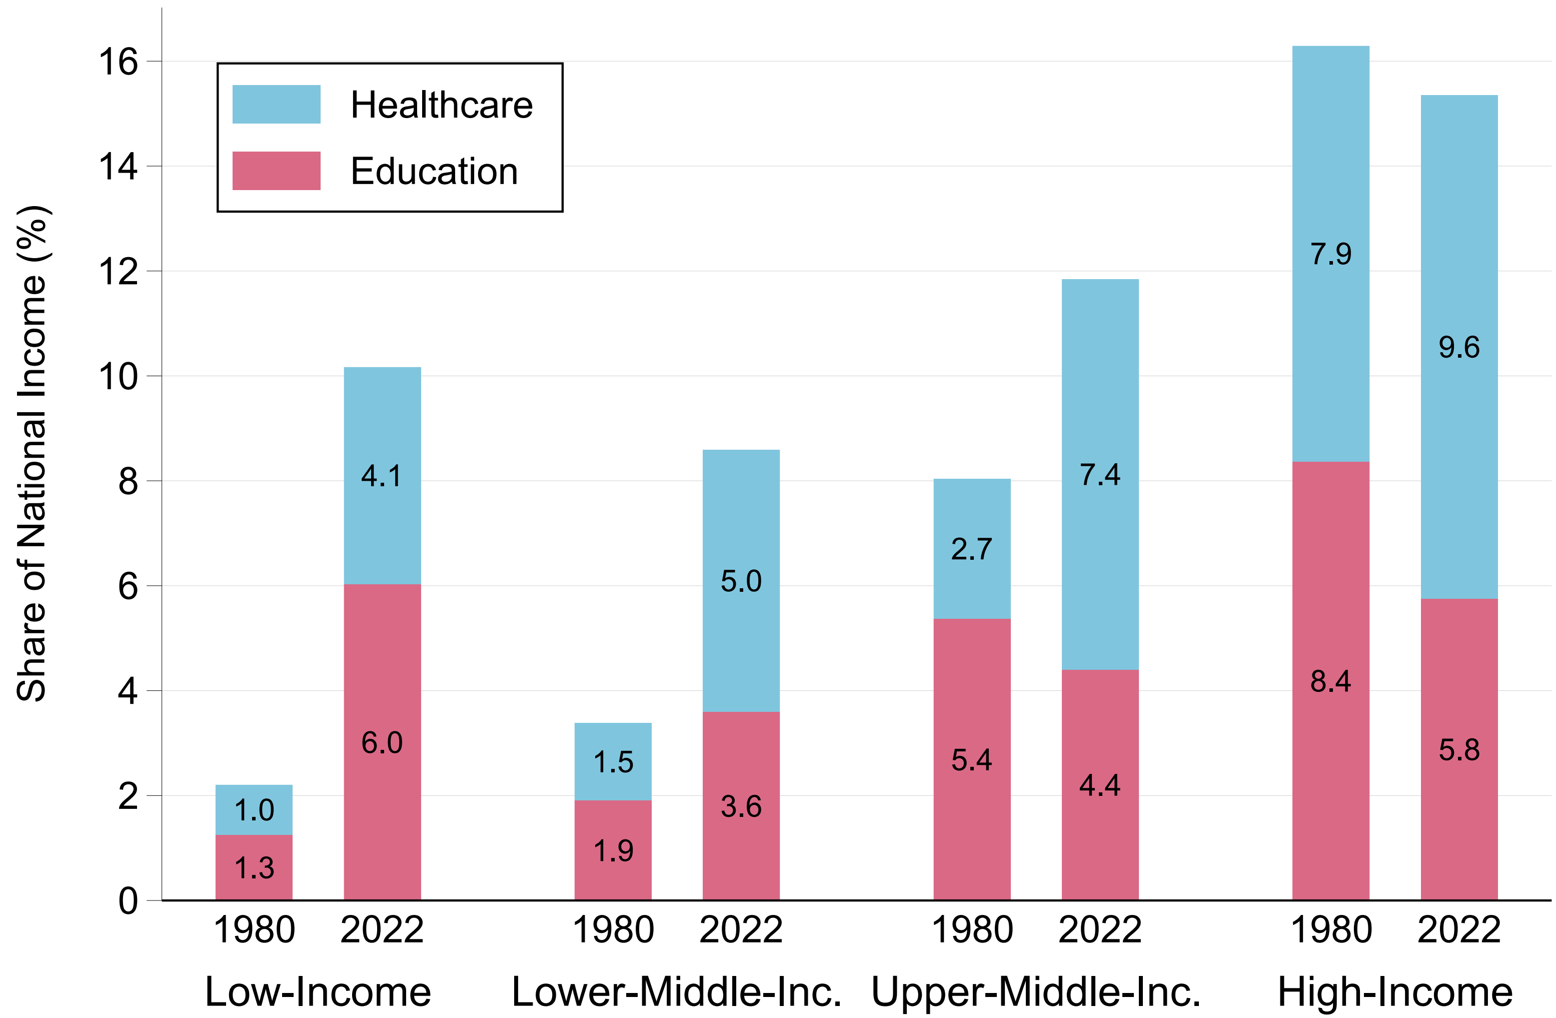

Notes. The figure plots the welfare value of public education and healthcare received by the bottom 50%, expressed as a share of national income, by country income group. The welfare value of education and healthcare is estimated using the future income specification, that is, accounting for future economic growth. Population-weighted averages across all countries in each group.

Figure A15 – Cost of Provision vs. Welfare Value of Public Services Received by the Bottom 50%

(a) Education

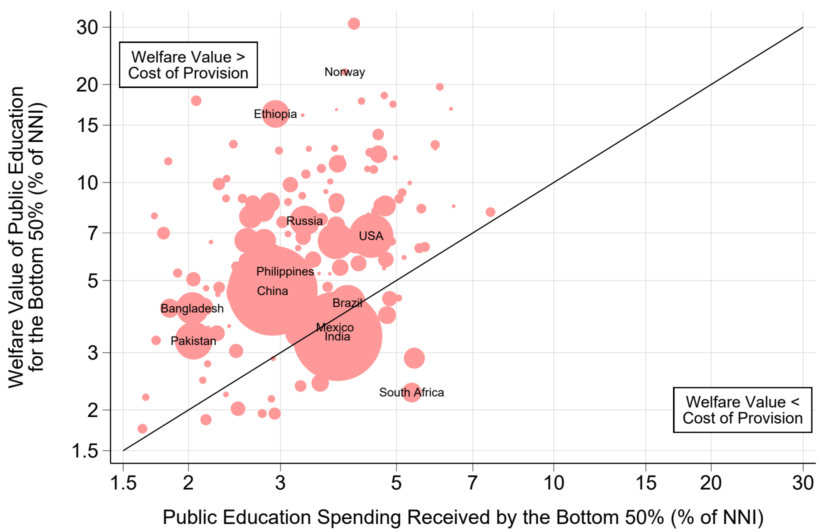

(b) Healthcare

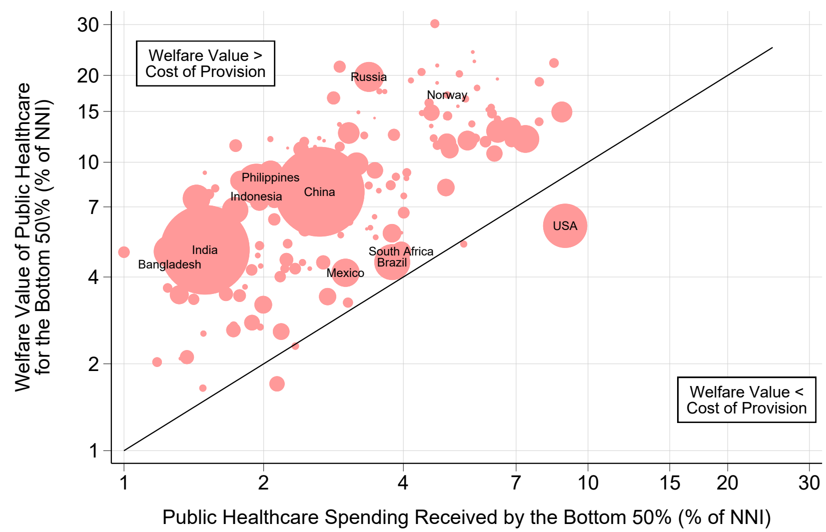

Notes. Panel A: the figure compares the public education transfer received by the bottom 50% in each country, expressed as a share of national income, depending on whether it is valued at cost of provision (x-axis) or at its net present value (y-axis). The net present value of education is calculated as the discounted expected monetary returns from an additional year of schooling. Panel B: the figure compares the public healthcare transfer received by the bottom 50% in each country, expressed as a share of national income, depending on whether it is valued at cost of provision (x-axis) or at its net present value (y-axis). The net present value of healthcare is calculated as the discounted expected monetary returns from living longer thanks to the healthcare system. Both panels: the net present value of education and healthcare is estimated using the future income approach, that is, accounting for future economic growth. The welfare value of public education or healthcare is higher than its cost of provision for countries above the 45-degree line; it is below cost of provision for countries below the 45-degree line.

Figure A16 – Total Public Education and Healthcare Transfers Received by the Global Bottom 20%: Cost of Provision versus Welfare Value

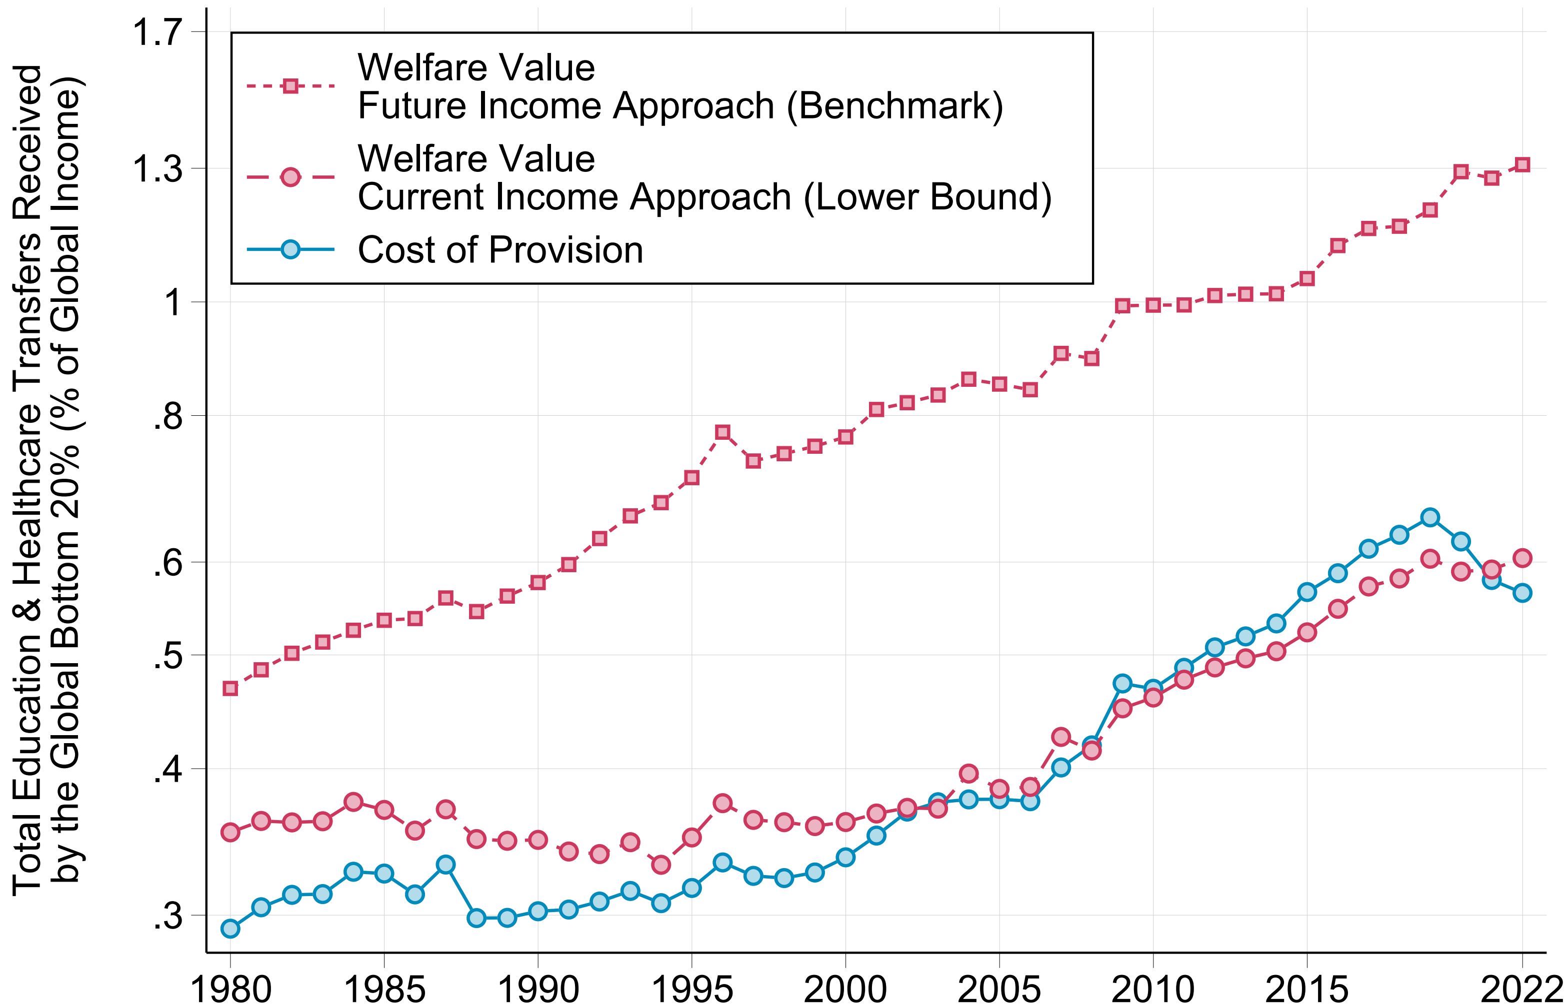

Notes. The figure compares the value of public education and healthcare received by the world's poorest 20%, expressed as a share of global income, depending on whether public services are valued at cost of provision or at their welfare value. Future income valuation: calculations of the net present value of education and healthcare account for future economic growth. Current income valuation: calculations of the net present value of education and healthcare assume zero future economic growth.

Figure A17 – Value of Healthcare Received by the Global Bottom 20% Under Alternative Specifications

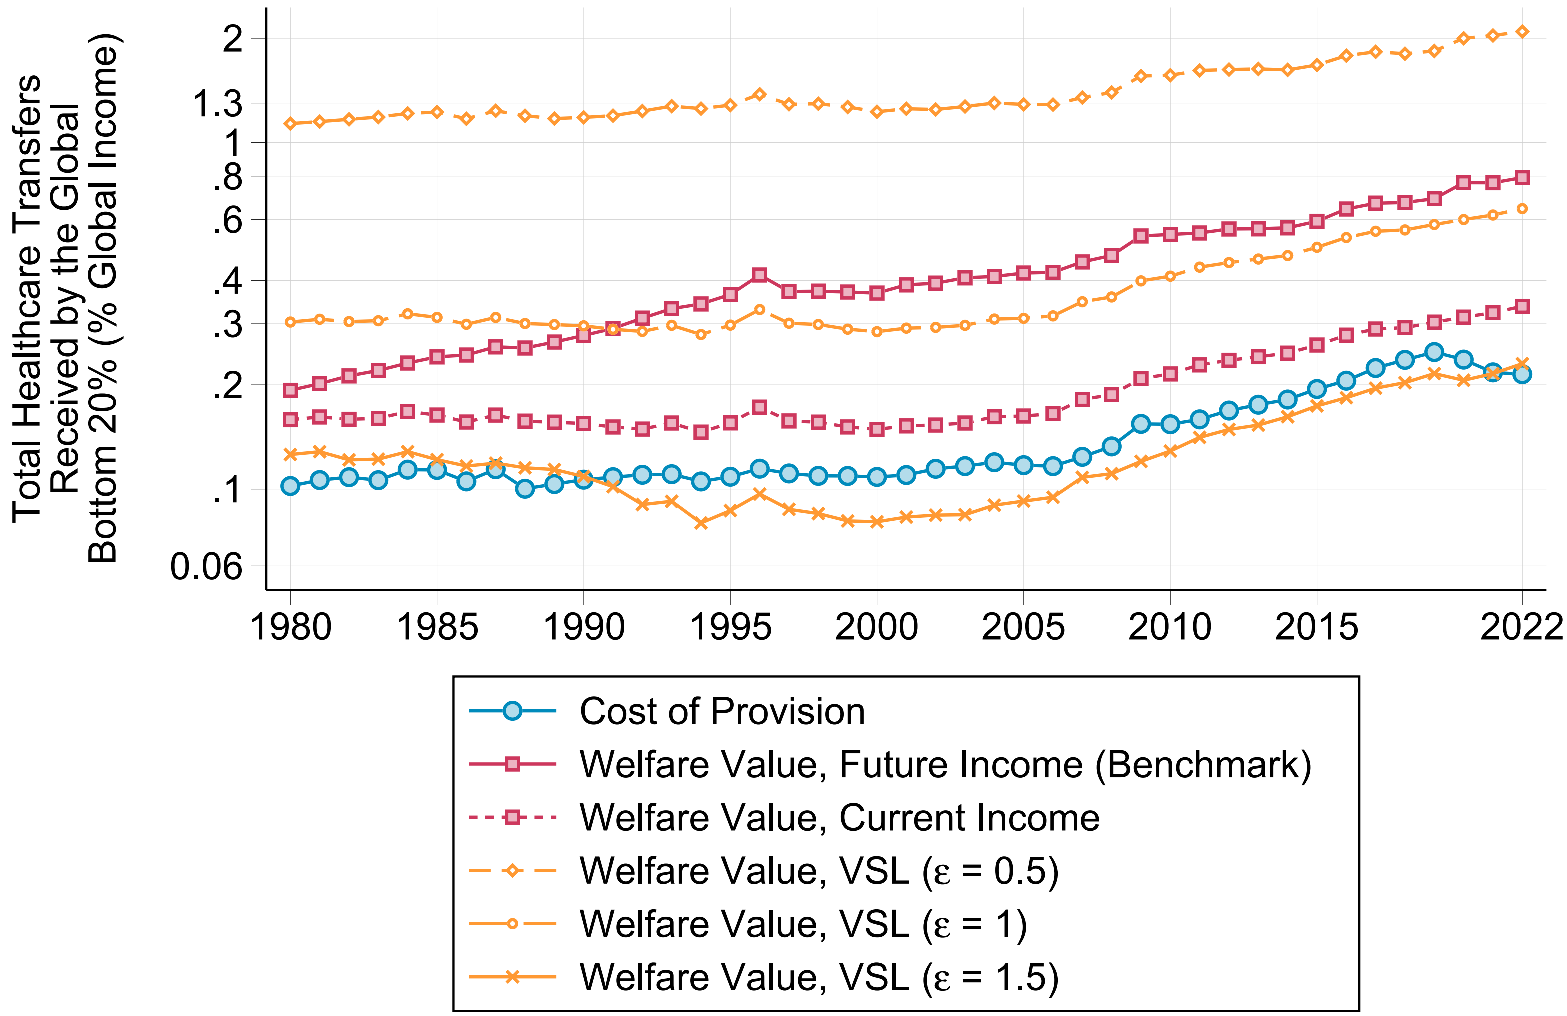

Notes. The figure compares the value of public healthcare received by the world's poorest 20%, expressed as a share of global income, depending on whether it is valued at cost of provision, at its welfare value using the future or current income specifications, or combining U.S. estimates of the value of a statistical life year with an income elasticity of VSL of ε.

Table A6 – Government Redistribution and Global Bottom 20% Growth: Sensitivity to Welfare Valuation of Public Services

	Global Bottom 20% Average Income (2021 PPP USD)		
	1980	2022	2022-1980
Pretax Income	326	587	80%
Posttax Income: Cost of Provision	508	1410	177%
Posttax Income: Welfare, Future Income	582	1899	227%
Posttax Income: Welfare, Current Income	524	1419	171%
Posttax Income: Welfare, VSL ($\varepsilon = 0.5$)	836	2615	213%
Posttax Income: Welfare, VSL ($\varepsilon = 1$)	545	1603	194%
Posttax Income: Welfare, VSL ($\varepsilon = 1.5$)	503	1318	162%

Notes. The table reports how results on the incidence of government redistribution on real income growth of the world's poorest 20% vary depending on assumptions regarding the valuation of public services. Cost of provision: value public education and healthcare at cost of provision. Welfare, current income: value public services based on discounted returns, estimated on current incomes. Welfare, future income: value public services based on discounted returns, estimated on projected future incomes. Welfare, VSL: value healthcare using estimates of a value of statistical life year, combining U.S. values with an income elasticity of VSL of ε.

Table A7 – Government Redistribution and Global Poverty Reduction: Sensitivity to Welfare Valuation of Public Services

	Global Poverty Headcount Ratio at $2.15 Per Day		
	1980	2022	2022-1980
Pretax Income	23.2%	14.5%	-38%
Posttax Income: Cost of Provision	19.0%	7.4%	-61%
Posttax Income: Welfare, Future Income	16.8%	4.4%	-74%
Posttax Income: Welfare, Current Income	18.6%	7.3%	-61%
Posttax Income: Welfare, VSL ($\varepsilon = 0.5$)	10.6%	1.6%	-85%
Posttax Income: Welfare, VSL ($\varepsilon = 1$)	18.0%	6.5%	-64%
Posttax Income: Welfare, VSL ($\varepsilon = 1.5$)	19.2%	7.6%	-60%

Notes. The table reports how results on the incidence of government redistribution on global poverty reduction vary depending on assumptions regarding the valuation of public services. Cost of provision: value public education and healthcare at cost of provision. Welfare, current income: value public services based on discounted returns, estimated on current incomes. Welfare, future income: value public services based on discounted returns, estimated on projected future incomes. Welfare, VSL: value healthcare using estimates of a value of statistical life year, combining U.S. values with an income elasticity of VSL of ε.

B.4. Data Sources

Table B8 – Data Coverage by Country

	Education		Healthcare		Social Protection	
	Spending	Progressivity	Spending	Progressivity	Spending	Progressivity
Afghanistan	1980-2022	2007-2011	1980-2022		1980-2022	2007
Albania	1980-2022	2002-2012	1980-2022	2016	1980-2022	2012
Algeria	1980-2022		1980-2022		1980-2022	
Angola	1980-2022	2000-2018	1980-2022		1980-2022	
Argentina	1980-2022	1980-2019	1980-2022	2017	1980-2022	2019
Armenia	1980-2022	1998-2019	1980-2022	2016	1980-2022	2018
Australia	1980-2022		1980-2022		1980-2022	
Austria	1980-2022	2003-2017	1980-2022	2003	1980-2022	2019
Azerbaijan	1980-2022	2002-2005	1980-2022	2016	1980-2022	2015
Bahrain	1980-2022		1980-2022		1980-2022	
Bangladesh	1980-2022	2000-2016	1980-2022	2003	1980-2022	2016
Belarus	1980-2022		1980-2022	2016	1980-2022	2019
Belgium	1980-2022	2003-2017	1980-2022	2003	1980-2022	2019

Belize	1980-2022		1980-2022		1980-2022	2009
Benin	1980-2022	2007-2018	1980-2022		1980-2022	
Bhutan	1980-2022	2003-2017	1980-2022		1980-2022	2007
Bolivia	1980-2022	1992-2019	1980-2022	2015	1980-2022	2017
Bosnia and Herzegovina	1980-2022	2001-2011	1980-2022	2016	1980-2022	2015
Botswana	1980-2022	2009-2015	1980-2022	2010	1980-2022	2015
Brazil	1980-2022	1981-2019	1980-2022	2003	1980-2022	2019
Brunei Darussalam	1980-2022		1980-2022			
Bulgaria	1980-2022	1995-2017	1980-2022	2016	1980-2022	2019
Burkina Faso	1980-2022	2009-2018	1980-2022	2003	1980-2022	2018
Burundi	1980-2022	1998-2013	1980-2022		1980-2022	
Cabo Verde	1980-2022	2007-2015	1980-2022		1980-2022	2007
Cambodia	1980-2022	1997-2008	1980-2022		1980-2022	
Cameroon	1980-2022	2007-2014	1980-2022		1980-2022	2014
Canada	1980-2022		1980-2022		1980-2022	
Central African Republic	1980-2022	2008-2008	1980-2022		1980-2022	
Chad	1980-2022	2011-2011	1980-2022	2003	1980-2022	2011
Chile	1980-2022	1990-2017	1980-2022	2013	1980-2022	2017
China	1980-2022	1988-2018	1980-2022	2003	1980-2022	2013

Colombia	1980-2022	1996-2019	1980-2022	2014	1980-2022	2019
Comoros	1980-2022	2004-2004	1980-2022	2003	1980-2022	2014
Costa Rica	1980-2022	1989-2019	1980-2022	2010	1980-2022	2019
Croatia	1980-2022	2009-2017	1980-2022	2016	1980-2022	2014
Cuba	1980-2022		1980-2022		1980-2022	
Cyprus	1980-2022	2004-2017	1980-2022	2016	1980-2022	2019
Czechia	1980-2022	2004-2017	1980-2022	2016	1980-2022	2019
Côte d’Ivoire	1980-2022	2002-2018	1980-2022	2003	1980-2022	2015
DR Congo	1980-2022	2004-2012	1980-2022	2012	1980-2022	2012
Denmark	1980-2022	2003-2017	1980-2022	2003	1980-2022	2019
Djibouti	1980-2022	1996-2017	1980-2022		1980-2022	2012
Dominican Republic	1980-2022	1996-2013	1980-2022	2003	1980-2022	2019
Ecuador	1980-2022	1994-2019	1980-2022	2003	1980-2022	2019
Egypt	1980-2022	1999-2019	1980-2022		1980-2022	2008
El Salvador	1980-2022	1991-2019	1980-2022	2017	1980-2022	2014
Equatorial Guinea	1980-2022		1980-2022		1980-2022	
Eritrea	1980-2022		1980-2022			
Estonia	1980-2022	2000-2017	1980-2022	2016	1980-2022	2019
Eswatini	1980-2022	2000-2016	1980-2022	2003	1980-2022	2016

Ethiopia	1980-2022	1995-2018	1980-2022	2003	1980-2022	2018
Finland	1980-2022	2003-2017	1980-2022	2003	1980-2022	2019
France	1980-2022	2003-2017	1980-2022	2003	1980-2022	2019
Gabon	1980-2022	2005-2017	1980-2022			2017
Gambia	1980-2022	1998-2015	1980-2022		1980-2022	2015
Georgia	1980-2022	2005-2019	1980-2022	2016	1980-2022	2018
Germany	1980-2022	2004-2017	1980-2022	2016	1980-2022	2019
Ghana	1980-2022	1987-2017	1980-2022	2003		2016
Greece	1980-2022	2003-2017	1980-2022	2016	1980-2022	2019
Guatemala	1980-2022	2000-2011	1980-2022	2003	1980-2022	2011
Guinea	1980-2022	2012-2018	1980-2022		1980-2022	2012
Guinea-Bissau	1980-2022	1993-2018	1980-2022		1980-2022	
Guyana	1980-2022	1992-1999	1980-2022		1980-2022	
Haiti	1980-2022	2001-2012	1980-2022		1980-2022	2001
Honduras	1980-2022	1991-2019	1980-2022	2011	1980-2022	2017
Hong Kong	1980-2022		1980-2022		1980-2022	
Hungary	1980-2022	1998-2017	1980-2022	2016	1980-2022	2019
Iceland	1980-2022	2003-2015	1980-2022		1980-2022	2019
India	1980-2022	1983-2017	1980-2022	2003	1980-2022	2011

Indonesia	1980-2022	1993-2019	1980-2022	2012	1980-2022	2019
Iran	1980-2022	2004-2019	1980-2022	2011	1980-2022	2011
Iraq	1980-2022	2006-2012	1980-2022		1980-2022	2012
Ireland	1980-2022	2004-2018	1980-2022	2003	1980-2022	2019
Israel	1980-2022		1980-2022	2003	1980-2022	
Italy	1980-2022	2003-2017	1980-2022	2016	1980-2022	2019
Jamaica	1980-2022	1990-2002	1980-2022		1980-2022	2017
Japan	1980-2022		1980-2022		1980-2022	
Jordan	1980-2022	2002-2013	1980-2022	2010	1980-2022	2010
Kazakhstan	1980-2022	2001-2010	1980-2022	2016	1980-2022	2017
Kenya	1980-2022	2005-2015	1980-2022	2003	1980-2022	2015
Kosovo	1980-2022	2009-2017	1980-2022	2016	1980-2022	2013
Kuwait	1980-2022		1980-2022		1980-2022	
Kyrgyzstan	1980-2022	2002-2003	1980-2022	2016	1980-2022	2013
Lao	1980-2022	2002-2018	1980-2022	2003	1980-2022	
Latvia	1980-2022	2004-2017	1980-2022	2016	1980-2022	2019
Lebanon	1980-2022	2004-2011	1980-2022		1980-2022	
Lesotho	1980-2022	2010-2017	1980-2022	2017	1980-2022	2017
Liberia	1980-2022	2007-2016	1980-2022		1980-2022	2016

Libya	1980-2022		1980-2022		1980-2022	
Lithuania	1980-2022	1998-2017	1980-2022	2016	1980-2022	2019
Luxembourg	1980-2022	2003-2017	1980-2022	2003	1980-2022	2019
Madagascar	1980-2022	2010-2012	1980-2022		1980-2022	
Malawi	1980-2022	2016-2019	1980-2022	2003	1980-2022	2016
Malaysia	1980-2022		1980-2022	2003	1980-2022	2016
Maldives	1980-2022	1998-2019	1980-2022		1980-2022	2016
Mali	1980-2022	1994-2018	1980-2022	2003	1980-2022	2014
Malta	1980-2022		1980-2022		1980-2022	2019
Mauritania	1980-2022	2000-2014	1980-2022	2003	1980-2022	2008
Mauritius	1980-2022	1999-2017	1980-2022	2003	1980-2022	2017
Mexico	1980-2022	1984-2018	1980-2022	2003	1980-2022	2018
Moldova	1980-2022	1998-2019	1980-2022	2016	1980-2022	2018
Mongolia	1980-2022	2002-2018	1980-2022	2016	1980-2022	2016
Montenegro	1980-2022	2007-2013	1980-2022	2016	1980-2022	2014
Morocco	1980-2022	1991-2013	1980-2022	2003	1980-2022	
Mozambique	1980-2022	2002-2014	1980-2022		1980-2022	
Myanmar	1980-2022	2015-2017	1980-2022	2003	1980-2022	2017
Namibia	1980-2022	2015-2015	1980-2022	2003	1980-2022	2015

Nepal	1980-2022	1995-2016	1980-2022	2003	1980-2022	2010
Netherlands	1980-2022	2004-2017	1980-2022	2003	1980-2022	2019
New Zealand	1980-2022		1980-2022		1980-2022	
Nicaragua	1980-2022	1993-2009	1980-2022	2009	1980-2022	2009
Niger	1980-2022	2005-2018	1980-2022		1980-2022	2014
Nigeria	1980-2022	1993-2018	1980-2022	2018	1980-2022	2018
North Macedonia	1980-2022	1996-2008	1980-2022	2016	1980-2022	
Norway	1980-2022	2003-2017	1980-2022	2003	1980-2022	2019
Oman	1980-2022		1980-2022		1980-2022	
Pakistan	1980-2022	1991-2018	1980-2022	2010	1980-2022	2010
Palestine	1980-2022	1996-2016				2016
Panama	1980-2022	1989-2019	1980-2022	2016	1980-2022	2019
Papua New Guinea	1980-2022	2009-2009	1980-2022		1980-2022	2009
Paraguay	1980-2022	1995-2019	1980-2022	2003	1980-2022	2019
Peru	1980-2022	1997-2019	1980-2022	2011	1980-2022	2019
Philippines	1980-2022	2003-2018	1980-2022	2003	1980-2022	2015
Poland	1980-2022	1998-2019	1980-2022	2016	1980-2022	2019
Portugal	1980-2022	2003-2017	1980-2022	2003	1980-2022	2019
Qatar	1980-2022		1980-2022		1980-2022	

Republic of the Congo	1980-2022	2011-2011	1980-2022	2003	1980-2022	2005
Romania	1980-2022	2000-2018	1980-2022	2016	1980-2022	2019
Russia	1980-2022	1994-2019	1980-2022	2019	1980-2022	2019
Rwanda	1980-2022	2010-2016	1980-2022		1980-2022	2013
Sao Tome and Principe	1980-2022	2010-2017	1980-2022		1980-2022	
Saudi Arabia	1980-2022		1980-2022		1980-2022	
Senegal	1980-2022	1995-2018	1980-2022	2003	1980-2022	2011
Serbia	1980-2022	2003-2019	1980-2022	2016	1980-2022	2015
Seychelles	1980-2022		1980-2022		1980-2022	
Sierra Leone	1980-2022	2003-2018	1980-2022		1980-2022	2018
Singapore	1980-2022		1980-2022		1980-2022	
Slovakia	1980-2022	2004-2017	1980-2022	2016	1980-2022	2019
Slovenia	1980-2022	2004-2017	1980-2022	2016	1980-2022	2019
Somalia	1980-2022		1980-2022		1980-2022	
South Africa	1980-2022	1994-2019	1980-2022	2003	1980-2022	2014
South Korea	1980-2022		1980-2022		1980-2022	
South Sudan	1980-2022	2009-2016	1980-2022			2009
Spain	1980-2022	2003-2017	1980-2022	2003	1980-2022	2019
Sri Lanka	1980-2022	2002-2019	1980-2022	2003	1980-2022	2016

Sudan	1980-2022	2009-2014	1980-2022		1980-2022	2009
Suriname	1980-2022		1980-2022		1980-2022	
Sweden	1980-2022	2003-2017	1980-2022	2003	1980-2022	2019
Switzerland	1980-2022	2008-2017	1980-2022		1980-2022	2019
Syrian Arab Republic	1980-2022	1997-2003	1980-2022			
Taiwan	1980-2022					
Tajikistan	1980-2022	1999-2015	1980-2022	2016	1980-2022	2011
Tanzania	1980-2022	1993-2018	1980-2022	2011	1980-2022	2014
Thailand	1980-2022	1988-2019	1980-2022		1980-2022	2019
Timor-Leste	1980-2022	2001-2014	1980-2022		1980-2022	2011
Togo	1980-2022	2006-2018	1980-2022	2015	1980-2022	2015
Trinidad and Tobago	1980-2022		1980-2022		1980-2022	
Tunisia	1980-2022	2005-2015	1980-2022	2003	1980-2022	2010
Turkey	1980-2022	2003-2019	1980-2022	2016	1980-2022	2019
Turkmenistan	1980-2022		1980-2022			
USA	1980-2022	1980-2019	1980-2022	1980-2022	1980-2022	1980-2022
Uganda	1980-2022	2010-2019	1980-2022	2016	1980-2022	2016
Ukraine	1980-2022	2002-2013	1980-2022	2016	1980-2022	2018
United Arab Emirates	1980-2022		1980-2022	2003	1980-2022	

United Kingdom	1980-2022	2005-2018	1980-2022	2003	1980-2022	2019
Uruguay	1980-2022	1989-2019	1980-2022	2003	1980-2022	2019
Uzbekistan	1980-2022	2000-2003	1980-2022	2016	1980-2022	2018
Venezuela	1980-2022	1989-2006	1980-2022	2013	1980-2022	2013
Vietnam	1980-2022	1992-2018	1980-2022	2003	1980-2022	2014
Yemen	1980-2022	1998-2014	1980-2022		1980-2022	2005
Zambia	1980-2022	2010-2015	1980-2022	2003	1980-2022	2015
Zimbabwe	1980-2022	2007-2019	1980-2022	2003	1980-2022	2019